

REPORT DOCUMENTATION PAGE			Form Approved OMB No. 0704-0188	
<small>Public reporting burden for this collection of information is estimated to average 1 hour per response, including the time for reviewing instructions, searching existing data sources, gathering and maintaining the data needed, and completing and reviewing the collection of information. Send comments regarding this burden estimate or any other aspect of this collection of information, including suggestions for reducing this burden, to Washington Headquarters Services, Directorate for Information Operations and Reports, 1215 Jefferson Davis Highway, Suite 1204, Arlington, VA 22202-4302, and to the Office of Management and Budget, Paperwork Reduction Project (0704-0188), Washington, DC 20503.</small>				
1. AGENCY USE ONLY (Leave blank)		2. REPORT DATE 14.Aug.02		3. REPORT TYPE AND DATES COVERED THESIS
4. TITLE AND SUBTITLE LARGE SCALE NETWORK DESIGN USING COMPOSIT VARIABLES: AN APPLICATION TO AIR MOBILITY COMMAND'S 30-DAY CHANNEL ROUTE NETWORK			5. FUNDING NUMBERS	
6. AUTHOR(S) 2D LT NIELSEN CHRISTOPHER A				
7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES) MASSACHUSETTS INSTITUTE OF TECHNOLOGY			8. PERFORMING ORGANIZATION REPORT NUMBER CI02-126	
9. SPONSORING/MONITORING AGENCY NAME(S) AND ADDRESS(ES) THE DEPARTMENT OF THE AIR FORCE AFIT/CIA, BLDG 125 2950 P STREET WPAFB OH 45433			10. SPONSORING/MONITORING AGENCY REPORT NUMBER	
11. SUPPLEMENTARY NOTES				
12a. DISTRIBUTION AVAILABILITY STATEMENT Unlimited distribution In Accordance With AFI 35-205/AFIT Sup 1			12b. DISTRIBUTION CODE	
13. ABSTRACT (Maximum 200 words)				
<div style="display: flex; justify-content: space-between; align-items: center;"> <div style="text-align: center;"> DISTRIBUTION STATEMENT A Approved for Public Release Distribution Unlimited </div> <div style="font-size: 2em; font-weight: bold;">20020829 030</div> </div>				
14. SUBJECT TERMS			15. NUMBER OF PAGES 196	
			16. PRICE CODE	
17. SECURITY CLASSIFICATION OF REPORT	18. SECURITY CLASSIFICATION OF THIS PAGE	19. SECURITY CLASSIFICATION OF ABSTRACT	20. LIMITATION OF ABSTRACT	

Large-Scale Network Design using Composite Variables: An Application to Air Mobility Command's 30-day Channel Route Network

by

Christopher A. Nielsen

Submitted to the Sloan School of Management
in Partial Fulfillment of the Requirements
for the Degree of Master of Science in Operations Research

ABSTRACT

Each month, the United States Air Force's Air Mobility Command is responsible for designing a large-scale air mobility network – called the *channel route network* – that is used to transport military personnel and cargo throughout the world. Traditionally, planning the channel route network has been a manual process, requiring numerous hours to generate a monthly channel route schedule. We formulate the monthly channel route scheduling problem using traditional network design formulations, and we use Dantzig-Wolfe decomposition to overcome tractability issues. The resulting linear programming bounds on the optimal integer solution are weak, however. Furthermore, the formulation is unable to capture complex operational rules and regulations that govern aircraft flows and cargo flows through the channel route network.

To overcome these challenges, we apply a variable redefinition technique known as *composite variable modeling*. Using this technique, we alleviate the need to explicitly include cargo flow decisions by implicitly capturing them in the design variables, which are then combined to create *composite variables*. A single composite variable represents the selection of one or more aircraft missions that completely cover a subset of cargo commodities. The resulting formulation is computationally superior to traditional network design formulations because it achieves tighter bounds – allowing excellent integer solutions to be found quickly – and it is able to capture complex operational rules and regulations. We illustrate that the nature of the formulation can be easily changed depending on the desired output by simply modifying the composite variable generation procedures, which does not compromise the structural properties of the formulation that allow it to achieve strong linear programming relaxations and short solution times. Furthermore, we develop the composite variable formulation to align with the functional description of the channel route planning process, allowing it to change with Air Mobility Command's evolving business practices.

Thesis Supervisor: Professor Cynthia Barnhart
Title: Professor of Civil and Environmental Engineering
Co-Director, Center for Transportation and Logistics

Technical Supervisor: Dr. Stephan E. Kolitz
Title: Principal Member of the Technical Staff,
The Charles Stark Draper Laboratory, Inc.

Technical Advisor: Professor Andrew P. Armacost
Title: Assistant Professor of Operations Research
Department of Management, US Air Force Academy, CO

**Large-Scale Network Design using Composite Variables: An Application to
Air Mobility Command's 30-day Channel Route Network**

by

Christopher A. Nielsen

B.S. Operations Research

United States Air Force Academy, 2000

SUBMITTED TO THE SLOAN SCHOOL OF MANAGEMENT IN PARTIAL
FULFILLMENT OF THE REQUIREMENTS OF THE DEGREE OF

MASTER OF SCIENCE IN OPERATIONS RESEARCH

at the

MASSACHUSETTS INSTITUTE OF TECHNOLOGY

June 2002

© 2002 Christopher A. Nielsen. All rights reserved.

The author hereby grants to MIT permission to reproduce and to distribute publicly paper and electronic
copies of this thesis document in whole or in part.

Signature of Author: Christopher A. Nielsen

Sloan School of Management

Interdepartmental Program in Operations Research

17 May, 2002

Approved by: Stephan E. Kolitz

Stephan E. Kolitz

The Charles Stark Draper Laboratory, Inc.

Technical Supervisor

Certified by: Cynthia Barnhart

Cynthia Barnhart

Professor, Civil and Environmental Engineering

Co-Director, Center for Transportation and Logistics

Thesis Advisor

Certified by: James B. Orlin

James B. Orlin

Edward Pennell Brooks Professor of Operations Research

Co-Director, Operations Research Center

[This Page Intentionally Left Blank]

Large-Scale Network Design using Composite Variables: An Application to Air Mobility Command's 30-day Channel Route Network

by

Christopher A. Nielsen

Submitted to the Sloan School of Management
on 17 May 2002, in Partial Fulfillment of the Requirements
for the Degree of Master of Science in Operations Research

ABSTRACT

Each month, the United States Air Force's Air Mobility Command is responsible for designing a large-scale air mobility network – called the *channel route network* – that is used to transport military personnel and cargo throughout the world. Traditionally, planning the channel route network has been a manual process, requiring numerous hours to generate a monthly channel route schedule. We formulate the monthly channel route scheduling problem using traditional network design formulations, and we use Dantzig-Wolfe decomposition to overcome tractability issues. The resulting linear programming bounds on the optimal integer solution are weak, however. Furthermore, the formulation is unable to capture complex operational rules and regulations that govern aircraft flows and cargo flows through the channel route network.

To overcome these challenges, we apply a variable redefinition technique known as *composite variable modeling*. Using this technique, we alleviate the need to explicitly include cargo flow decisions by implicitly capturing them in the design variables, which are then combined to create *composite variables*. A single composite variable represents the selection of one or more aircraft missions that completely cover a subset of cargo commodities. The resulting formulation is computationally superior to traditional network design formulations because it achieves tighter bounds – allowing excellent integer solutions to be found quickly – and it is able to capture complex operational rules and regulations. We illustrate that the nature of the formulation can be easily changed depending on the desired output by simply modifying the composite variable generation procedures, which does not compromise the structural properties of the formulation that allow it to achieve strong linear programming relaxations and short solution times. Furthermore, we develop the composite variable formulation to align with the functional description of the channel route planning process, allowing it to change with Air Mobility Command's evolving business practices.

Thesis Supervisor:	Professor Cynthia Barnhart
Title:	Professor of Civil and Environmental Engineering Co-Director, Center for Transportation and Logistics
Technical Supervisor:	Dr. Stephan E. Kolitz
Title:	Principal Member of the Technical Staff, The Charles Stark Draper Laboratory, Inc.
Technical Advisor:	Professor Andrew P. Armacost
Title:	Assistant Professor of Operations Research Department of Management, US Air Force Academy, CO

[This Page Intentionally Left Blank]

**THE VIEWS EXPRESSED IN THIS ARTICLE
ARE THOSE OF THE AUTHOR AND DO NOT
REFLECT THE OFFICIAL POLICY OR
POSITION OF THE UNITED STATES,
DEPARTMENT OF DEFENSE, OR THE U.S.
GOVERNMENT**

Table of Contents

1	INTRODUCTION	15
1.1	Research Scope – Channel Routes	17
1.2	Overview of Thesis	19
2	AIR MOBILITY COMMAND AND THE CURRENT CHANNEL ROUTE PLANNING AND EXECUTION PROCESS.....	23
2.1	Air Mobility Command	23
2.1.1	<i>AMC Customers</i>	<i>24</i>
2.1.2	<i>The Tanker Airlift Control Center</i>	<i>25</i>
2.1.3	<i>AMC's Three Objectives</i>	<i>26</i>
2.2	The Physical System	34
2.2.1	<i>Cargo Items.....</i>	<i>35</i>
2.2.2	<i>Aerial Ports</i>	<i>38</i>
2.2.3	<i>Aircraft Assets.....</i>	<i>40</i>
2.2.4	<i>Aircrews</i>	<i>45</i>
2.3	The Current Planning & Execution Process	46
2.3.1	<i>Literature Review – Historical Models at AMC.....</i>	<i>47</i>
2.3.2	<i>Channel Route Planning Tools</i>	<i>49</i>
2.3.3	<i>The Planning Process.....</i>	<i>50</i>
2.4	Motivation for Models	62
2.4.1	<i>Readiness.....</i>	<i>62</i>
2.4.2	<i>Customer Service</i>	<i>62</i>
2.4.3	<i>Net Operating Result.....</i>	<i>63</i>
2.5	Summary	63
3	FUNCTIONAL ANALYSIS & TRADITIONAL MODELING APPROACHES	65
3.1	Functional View of the Channel Route Scheduling Problem	66

3.2	An Improved Operational View	66
3.3	Static Graph Representation.....	69
3.4	Time-Space Graph Representation.....	70
3.4.1	<i>General Time-Space Graph</i>	70
3.4.2	<i>Incorporating Constraints in the Time-Space Graph</i>	73
3.5	Problem Classification.....	76
3.5.1	<i>Multicommodity Network Flow Formulations</i>	77
3.5.2	<i>Network Design Formulation</i>	79
3.6	Network Design Literature Review	80
3.7	Channel Route Formulations	81
3.7.1	<i>Commodity Definition</i>	81
3.7.2	<i>Channel Route Arc-Flow Formulation</i>	82
3.7.3	<i>Channel Route Path-Flow Formulation</i>	88
3.8	Decomposition	92
3.8.1	<i>Relaxed Constraints</i>	94
3.8.2	<i>Master Level Scheduler</i>	95
3.8.3	<i>Shortest Path Subproblems</i>	96
3.8.4	<i>Implementation</i>	100
3.9	Results & Analysis	103
3.10	Summary	106
4	COMPOSITE VARIABLE FORMULATION	109
4.1	Motivation.....	109
4.1.1	<i>Composite Variable Overview</i>	110
4.2	Literature Review.....	115
4.3	Formulation.....	117
4.4	Composite Variable Generation	121
4.4.1	<i>Step 1: Flight Sequences</i>	121
4.4.2	<i>Step 2: Generic Missions</i>	123
4.4.3	<i>Step 3: Single-route Composite Missions</i>	125
4.4.4	<i>Step 4: Multi-route Composite Missions</i>	131

4.5	Composite Variable Generation Considerations.....	133
4.6	Results & Analysis	138
4.7	Summary	139
5	RESULTS AND ANALYSIS.....	141
5.1	Metrics	142
5.2	Objective Functions	143
5.3	Parameters	144
5.4	Computational Performance	145
5.4.1	<i>Objective Functions.....</i>	<i>145</i>
5.4.2	<i>Parameters.....</i>	<i>146</i>
5.5	Solution Analysis.....	152
5.5.1	<i>Objective Functions.....</i>	<i>152</i>
5.5.2	<i>Parameters.....</i>	<i>155</i>
5.6	Summary	162
6	TOWARDS A DECISION SUPPORT TOOL	165
6.1	What is a Decision Support System.....	166
6.2	A Decision Support Tool within the Current Planning Process	166
6.3	Desired Capabilities	168
6.3.1	<i>Initial Cut Decision Support.....</i>	<i>168</i>
6.3.2	<i>Execution Month Decision Support.....</i>	<i>169</i>
6.3.3	<i>Strategic Decision Support</i>	<i>170</i>
6.4	Enhancing CRM-C	171
6.4.1	<i>Hierarchical Decomposition.....</i>	<i>171</i>
6.4.2	<i>Extreme Route Models and Column Generation.....</i>	<i>172</i>
6.4.3	<i>Additional Composite Variables</i>	<i>172</i>
6.5	Summary	175
7	SUMMARY AND FUTURE RESEARCH BEYOND CHANNEL ROUTES.....	177

7.1	Thesis Summary.....	177
7.2	Future Research	180
APPENDIX A: FORMULATIONS		185
APPENDIX B: GLOSSARY OF ACRONYMS		189
REFERENCES		191

List of Figures

Figure 1-1: SDMI process improvement committees	16
Figure 1-2: USTRANSCOM organizational structure	17
Figure 2-1: Percentage of cargo shipped by AMC customers in FY99 and FY00	25
Figure 2-2: TACC organizational chart	26
Figure 2-3: Mission areas by percentage of total flying hours for FY99 and FY00	28
Figure 2-4: Channel cargo shipped by type for FY99 and FY00	35
Figure 2-5: Standard pallet and cargo loaders	36
Figure 2-6: Cargo flows originating at (a) CONUS APOE, or (b) OCONUS APOE	37
Figure 2-7: CONUS APOEs and associated regions throughout the world	39
Figure 2-8: OCONUS hubs in the channel route network	40
Figure 2-9: The channel route planning period and execution month	50
Figure 2-10: The channel route planning timeline of events	51
Figure 2-11: The current channel route scheduling process	53
Figure 3-1: Functional View of the Channel Route Planning Process	66
Figure 3-2: The airline schedule planning problem	67
Figure 3-3: An operational view of (a) the current planning process and (b) a planning process that simultaneously considers all relevant inputs	68
Figure 3-4: Graph Representations of the Channel Route Network	70
Figure 3-5: Time-Space Graph of the Channel Route Network	71
Figure 3-6: Aircraft and cargo flow through the time-space graph	72
Figure 3-7: Construction of a time-space graph for a notional channel route network	74
Figure 3-8: Network design choices	79
Figure 3-9: Aircraft Routes and Cargo Paths as columns in CRM-P	91
Figure 3-10: Structure of columns in CRM-C	91
Figure 3-11: Dantzig-Wolfe Decomposition applied to CRM-P	93
Figure 3-12: Label updating procedure for multilabel shortest path algorithm	100
Figure 3-13: Multi-Label Shortest Path Example	101
Figure 3-14: Initial cargo path and aircraft route columns	103
Figure 3-15: Column Generation Flow Chart	104
Figure 3-16: Example cargo path column from decomposition solution	107
Figure 4-1: Channel route network with two aerial ports and two aircraft types	111
Figure 4-2: Removing fractional solutions using coefficient reduction	112

Figure 4-3: A composite variable constructed of different aircraft types and missions	113
Figure 4-4: Covering constraints using composite variables	114
Figure 4-5: Commercial composite example.....	120
Figure 4-6: Composite Variable Generation Steps	121
Figure 4-7: Generic Composite Mission feasibility check procedure	124
Figure 4-8: Example of a generic mission.....	126
Figure 4-9: Network for Single-route Composite Mission – Case 2.....	128
Figure 4-10: Network for Single-route Composite Mission with multiple cargo loadings	129
Figure 4-11: Example Network for Single-route Composite Mission generation	130
Figure 4-12: Example Network for Multi-route Composite Mission	131
Figure 4-13: Notional range-payload curve	133
Figure 4-14: Example Network for Multi-route Composite Mission generation	134
Figure 4-15: Influence of StartFreq on aircraft balance of flow	135
Figure 5-1: Impact of (a) PHT and (b) Fract on the number of columns in CRM-C.....	147
Figure 5-2: Structure of flight sequences in Seq Base, Seq B, and Seq E	149
Figure 5-3: Impact of flight sequences on the number of columns in CRM-C	149
Figure 5-4: Impact of StartFreq on the number of columns in CRM-C	150
Figure 5-5: Impact of Multiple Parameters on the number of columns in CRM-C.....	152
Figure 5-6: Metric values vs. Objective Function	153
Figure 5-7: Metric values vs. Port Hold Time (PHT)	156
Figure 5-8: Metric values vs. Fraction of Available Aircraft Capacity	157
Figure 5-9: Metric values vs. Flight Sequences	158
Figure 5-10: Metric values vs. StartFreq	159
Figure 5-11: Effect of simultaneously varying parameters on metrics	161
Figure 6-1: Planning (a) without and (b) with an optimization-based decision support tool....	167
Figure 7-1: Proposed Barrelmaster Decomposition	181
Figure 7-2: A composite variable for the logistics pipeline problem	183

List of Tables

Table 2-1: Priority, predictability, and percent of total flying hours by mission area	30
Table 2-2: Organic aircraft characteristics.....	41
Table 2-3: FY00 Civil Reserve Air Fleet (CRAF) Participants	43
Table 2-4: Aircrew characteristics	46
Table 3-1: Example Channel Route Network Characteristics	86
Table 3-2: Size of CRM-A for various Channel Route Network instances.....	87
Table 3-3: Number of constraints in CRM-P for various Channel Route Network instances	92
Table 3-4: Initial Cargo Paths Generated from an Initial Channel Route Mission	102
Table 3-5: East Data Set Characteristics.....	105
Table 4-1: Single-Route Composite Missions	128
Table 4-2: Single-route Composite Missions with multiple cargo loadings	130
Table 4-3: Multi-Route Composite Missions	131
Table 4-4: Data set characteristics	138
Table 4-5: Computational Results of baseline CRM-C solution	139
Table 5-1: Metrics used to analyze CRM-C.....	142
Table 5-2: Objective functions tested	143
Table 5-3: Number of aircraft available of each aircraft type.....	144
Table 5-4: Summary of parameters tested	144
Table 5-5: Influence of objective functions on the computational performance of CRM-C	146
Table 5-6: Summary of parameters' effect on the size of CRM-C	151
Table 5-7: Summary of objective function performance	154
Table 5-8: Summary of parameters' effect on the solution of CRM-C.....	160
Table 5-9: Size and solution time with Fract = 0.10, Seq B, and Cost objective	162

[This Page Intentionally Left Blank]

1 Introduction

The United States military logistics system is crucial for success during contingency operations and times of conflict. However, equally important is the posture of the military logistics system during peacetime. A healthy peacetime logistics system ensures that military personnel around the world continually have the resources required to maintain the high level of readiness required to respond at a moments notice.

The amount of cargo and the number of personnel moved during the Gulf War is impressive. Yet, most analyses of this logistical feat have concluded that success was a result of innovative, ad-hoc solutions that completely bypassed the established logistical system [46]. Continually patching the logistics system each time the military commits to a contingency does not allow the military to implement its wartime doctrine and project its power fully. This is especially true in the current state of the world, which does not allow for the large planning lead times that existed prior to the Gulf War. Consequently, the peacetime logistics system must have the ability to be pushed to the limits with minimal hiccups when the U.S. military commits its forces.

Attaining a superior peacetime logistics system is a daunting task. The logistics system is a large-scale, complex system that involves many detailed, interrelated decisions, processes, and organizations. To attain a superior logistics system the military must ensure that the latest planning and scheduling technologies are being exploited for high-level strategic analysis of the logistics system, as well as its day-to-day operation.

A critical component of the logistics system is the *Defense Transportation System (DTS)*. The *United States Transportation Command (USTRANSCOM)* manages all *Department of Defense (DoD)* transportation, which occurs over the DTS. This includes air, land, and sea modes of transportation. Since the end of the Cold War, the military logistics system has undergone major changes at the hand of a concept called *focused logistics*. Focused logistics is “the ability to provide the joint force the right personnel, equipment, and supplies in the right place, at the right time, and in the right quantity, across the full range of military operations” [37]. This concept is in contrast to the logistics system that prevailed during the Cold War in which large inventories of required and just-in-case materiel were stockpiled at numerous locations around the world. The current military force is smaller, is more constrained financially, and does not enjoy nearly as many permanent locations throughout the world. In this environment, materiel cannot be stockpiled in large quantities and must, in fact, flow through the logistics system continually and have the ability to seamlessly transition between various modes of transportation.

The *Strategic Distribution Management Initiative (SDMI)* was implemented by USTRANSCOM to streamline the entire military logistics system. This initiative is aimed at accomplishing USTRANSCOM’s strategic vision of providing “timely, customer-focused global mobility in peace and war through efficient, effective, and integrated transportation from origin to destination” [52]. The purpose of SDMI is to analyze and refine the processes within the logistics system with a goal of providing superior end-to-end logistical support for the military warfighter. SDMI is composed of four committees (see *Figure 1-1* [52]), each focused on improving a piece of the military logistics system.




Figure 1-1: SDMI process improvement committees

Transportation processes are vital to the success of SDMI. The architects of SDMI must find ways to integrate all three modes of transportation to ensure that cargo and personnel continuously flow through the military logistics system. This requires high-level analysis to develop guiding strategic concepts. It is also critical that the implementers of the strategic concepts – the day-to-day operational planners and schedulers – use scheduling methods and techniques that will ensure that the strategic concepts permeate the entire military logistics system.

1.1 Research Scope – Channel Routes

An integral piece of the DTS is air transportation, or *air mobility*. Air mobility allows the military to move materiel quickly over large distances. The United States Air Mobility Command (AMC), one of three *transportation component commands* (TCCs) within USTRANSCOM, is responsible for air transportation (see *Figure 1-2* [51]).

To fulfill its mission, AMC must manage a large, complex air mobility network. In fiscal year 2000 (FY00), AMC planned and executed, on average, 300 missions each day. This equates to approximately 1000 different daily flight legs, which took place in 156 of the world's 190 countries [48]. The throughput required is enormous; in FY00 AMC moved 163,174 tons of cargo, airlifted more than 31,350 medical patients and 338,372 military personnel worldwide, and flew 133,297 hours [52]. To ensure that the air transportation network remains capable of this level of throughput and performance, AMC is engaged in all levels of analysis, ranging from high-level, strategic analysis to day-to-day operational level planning and execution. Examples

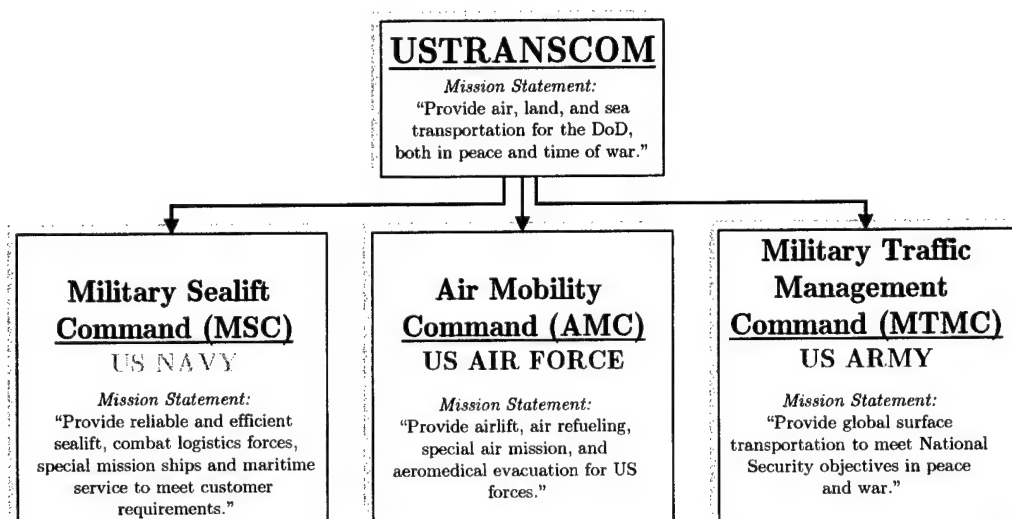


Figure 1-2: USTRANSCOM organizational structure

of high-level strategic analysis include new aircraft acquisitions and continual analysis of what-if wartime scenarios. The day-to-day operational level planning involves detailed schedule generation for each aircraft mission.

A key component of the air mobility network is the *channel route network*. This network is used to move cargo and personnel from the Continental United States (CONUS) to locations throughout the world on a *predictable, recurring* basis. The channel route network is used to *distribute* materiel and *sustain* military personnel located throughout the world during peacetime. The channel route network is complex, requiring intricate coordination with active duty Air Force flying squadrons, Air Force Reserve and Air National Guard flying squadrons, and contracted commercial air carriers. Although the channel route network is a major component of AMC's peacetime operations, there are other mission areas – intra-service and multinational training exercises, special assignment airlift missions (SAAMs), ongoing contingency operations – that require air mobility. Among these, missions occurring over the channel route network often have the largest planning lead-time, but the lowest priority. In contrast, other mission areas, such as SAAMs, have a high priority but are short notice. Complicating matters further is that a common set of aircraft are used to service each mission area. In other words, there is not a set of aircraft allocated solely to operate channel routes. This often leaves the schedule of channel route missions highly susceptible to changes and cancellations. When a channel route mission is canceled, cargo in the channel route network becomes either stagnant, increasing customer wait time and causing delays to ripple through the logistics system, or AMC must pay a commercial carrier a premium rate to move the cargo.

In addition to the cargo that is transported over the channel route network, other considerations influence the design and operation of the channel route network. For instance, flying hour requirements (which are used as a proxy for aircrew training and readiness) are paramount among AMC's overall goals and objectives. During the 1990's AMC relied on the Gulf War and other contingencies (e.g., Bosnia) as the major source of flying hours. Now that those conflicts have ended, the channel route network has been the main source of flying hours. Consistently meeting the flying hour requirements set forth by the United States Congress is difficult, however, as a result of the channel route schedule's propensity to change.

Currently, planning and scheduling channel route missions is a labor intensive, manual process that undergoes many time-consuming iterations before an implementable schedule is created. To deal with the overwhelming amount of information that must be incorporated into the channel route schedule, the scheduling process has traditionally disaggregated many highly related decisions. Consequently, the system-wide interactions and impact of tightly coupled planning decisions can only be evaluated on a limited scope, relying on planners' intuition and

experience. In a system of this magnitude, intuition and experience do not always yield the optimal use of resources.

The focus of this thesis is to improve the portion of the *channel route network* that is used to transport *cargo*. Although the channel route network associated with cargo movements is only a subset of the entire logistics system, it is expensive to operate and it is subject to many changes. Thus, even small improvements can yield big returns in terms of operating costs, resources consumed, and the ability to dynamically respond to system changes. The goal of this research is to explore the applicability of recent modeling technologies that have been used to model similar large-scale systems such as United Parcel Service's (UPS) Next Day Air delivery network [3],[4]. Specifically, the models developed in this thesis are aimed at improving the monthly channel route schedule.

Three key ideas motivate the research in this thesis. First, manually creating and maintaining a channel route schedule is time consuming, leaving little or no time to analyze alternate channel route network structures. An automated approach can drastically reduce the time and effort required to generate a channel route schedule allowing channel route planners to evaluate the system-wide effects of various solutions. Second, an optimization-based approach can simultaneously – rather than sequentially – consider all relevant inputs. Consequently, highly related decisions remain coupled during the scheduling process. Finally, we seek a method that improves the performance of the channel route network and that has the ability to account for AMC's three objectives (readiness, customer service, and net operating result).

1.2 Overview of Thesis

In this thesis we provide an overview of the current channel route planning and execution process, we present the advantages and disadvantages of modeling the channel route scheduling problem using traditional network design formulations, and we introduce a composite variable formulation that improves upon traditional network design formulations of the channel route network. The remaining chapters are organized as follows:

CHAPTER 2: *Air Mobility Command and the Current Channel Route Planning and Execution System*

The purpose of this chapter is to introduce the channel route network in the context of all of AMC's mission areas. We present a detailed description of the physical channel route network and the current scheduling and execution process. Furthermore, we

highlight shortcomings of the current planning process, which we use to motivate the models and formulations developed in the following chapters.

CHAPTER 3: *Functional Analysis and Traditional Modeling Approaches*

This chapter presents a functional analysis of the channel route scheduling problem. Furthermore, we formulate the problem using a traditional network design formulation in which cargo flows and aircraft flows are modeled using path variables. We develop and implement a decomposition approach based on Dantzig-Wolfe decomposition that utilizes shortest path subproblems to generate cargo flow and aircraft flow variables on the fly. Computational results are presented and motivation for the composite variable formulation is discussed.

CHAPTER 4: *Composite Variable Formulation*

In this chapter, we model the channel route scheduling problem using a composite variable formulation, which has superior computational behavior compared to traditional network design formulations of the channel route network. We implement the composite variable formulation and compare the computational results to those obtained in Chapter 3. Specifically, we use the computational results to illustrate the improved quality of the composite variable formulation's linear programming (LP) relaxation. The improved LP-relaxation results in faster solution times and allows us to significantly increase the problem size. We also illustrate the ability of the composite variable formulation to capture operational nuances that are difficult to model using traditional network design formulations.

CHAPTER 5: *Results and Analysis*

In this chapter, we capitalize on the superior computational behavior of the composite variable formulation to investigate various objective functions and different values of the input parameters to the composite variable generation procedures. The results and analysis in this chapter serves two purposes. First, we present an overall picture of how the size and tractability of the composite variable formulation, as well as the solution quality, change under different objective functions and parameter values. Second, the range of analysis illustrates the flexibility provided by composite variable modeling. Specifically, we demonstrate that the composite variable formulation can easily be manipulated through the composite variable generation procedures without destroying

the structure of the formulation that yields strong LP-relaxations and short solution times.

CHAPTER 6: *Towards a Decision Support Tool*

In this chapter we discuss the transition of the composite variable formulation into an optimization-based decision support tool, which has two potential benefits. One, it would allow channel route planners to make more informed decisions at each step of the current planning process. Two, it would reduce the time required to generate a channel route schedule and allow various solutions to be tested and evaluated. We discuss the impact of such a tool on the current planning process. Equally important, we highlight the ability of the research in this thesis – which we develop to meet the functional description of the channel route planning process – to change with evolving business practices. We also highlight desired capabilities of an optimization-based decision support tool. Furthermore, we outline specific enhancements that would allow the composite variable formulation we develop in this thesis to perform the desired capabilities.

CHAPTER 7: *Summary and Future Research*

This chapter summarizes the work presented in the thesis and discusses several areas of future research.

[This page intentionally left blank]

2 Air Mobility Command and the Current Channel Route Planning and Execution Process

Chapter 1 provided an introduction to channel routes and AMC in the context of the entire military logistics system. This chapter has three objectives. The first is to explore channel routes in the context of all AMC operations and AMC's three broad operational objectives of *readiness*, *customer service*, and *net operating result (NOR)*. The second is to introduce the physical components of the channel route network, which are the *aircraft*, *aerial ports*, and *cargo*. Finally, the third is to describe the current planning and execution process used to schedule and execute channel routes.

2.1 Air Mobility Command

Whenever called upon, AMC must be able to support a national war effort. This readiness requirement is at the crux of AMC's peacetime existence. Specifically, AMC must ensure that Air Force pilots maintain an adequate level of training and readiness during peacetime so that they are prepared to go to war at a moment's notice. To do this, AMC must

schedule peacetime flying operations to meet a predetermined annual readiness requirement defined by the *Flying Hour Program* (FHP). On a day-to-day basis, the organization responsible for scheduling all aircraft missions is the *Tanker Airlift Control Center* (TACC). The TACC was created in the early 1990's to provide "one-stop shopping" for all DoD customers seeking air mobility support [49].

Although pilot readiness is the number one priority, there are two other objectives AMC strives to fulfill. These two objectives are to provide a predictable level of service for AMC customers and to recover as much of the operating costs (measured by NOR) as possible.


The purpose of this section is to describe AMC's customers, introduce the role of the TACC, and outline AMC's three broad operational objectives.

2.1.1 AMC Customers

The Department of Defense (DoD) is composed of nine major commands, each with a Commander in Chief (CINC). Of the nine major commands, five are defined by geographic boundaries. The CINCs of these geographic major commands are responsible for all operations within the Command's geographic boundaries. The remaining four major commands are not defined by geographic regions, but rather by functional capabilities. The CINCs of the functional major commands have responsibility for these functions worldwide. The U.S. Transportation Command (USTRANSCOM) is one of the four functional major commands. USTRANSCOM's mission is "to provide air, land, and sea transportation for the Department of Defense, both in time of peace and time of war." [52]. USTRANSCOM is ultimately responsible for all transportation that occurs over the Defense Transportation Network (DTS), defined to be the "portion of the nation's transportation infrastructure which supports Department of Defense common-user-transportation needs across the range of military operations. It consists of those common-user military and commercial assets, services, and systems organic to, contracted for, and controlled by the Department of Defense" [52].

The DTS is modeled as a "customer-seller relationship between the provider (i.e., USTRANSCOM) and the customer (i.e., CINCs and services)" [52]. Shipments moving within the DTS result from requirements specified by authorized customers, which include the CINCs of the five geographic major commands, all military services (e.g., Air Force, Army, Navy), the Department of State (DOS), the Defense Logistics Agency (DLA), the Defense Commissary Agency (DeCA), the Military Postal System, and other federal agencies such as the Federal Bureau of Investigation (FBI) [51].

**FY99 AMC Customers by Short Tons
Shipped via Channel Routes**



**FY00 AMC Customers by Short Tons
Shipped via Channel Routes**




Figure 2-1: Percentage of cargo shipped by AMC customers in FY99 and FY00

The largest AMC customer has typically been the Air Force, followed closely by the Army, Navy, and the DLA, as shown in *Figure 2-1*.

2.1.2 The Tanker Airlift Control Center

At the end of the Cold War the *Military Airlift Command* (MAC) controlled the Air Force's *strategic airlifters* (i.e., large cargo carrying aircraft) and the *Strategic Air Command* (SAC) controlled the long-range bombers (e.g., Boeing B-52) and all air-refueling aircraft. The *Tactical Air Command* controlled fighter aircraft (e.g., General Dynamics F-16). This changed in June 1992 when AMC and *Air Combat Command* (ACC) were formed. AMC assumed responsibility of the strategic airlift, and air refueling assets and ACC assumed responsibility of the Air Force's fighter and bomber aircraft [49].

An organization formed within AMC was the *Tanker Airlift Control Center* (TACC), which is co-located with AMC Headquarters at Scott Air Force Base, Illinois. The TACC provides centralized scheduling and execution of all AMC strategic airlift and air refueling assets [49]. The need to simplify and streamline access to airlift assets for all customers seeking military airlift motivated the creation of the TACC [50].

Ten *directorates* exist within the TACC (see *Figure 2-2*). All AMC missions are categorized into one of five different mission areas (described in §2.1.3). Several of the directorates handle the planning and scheduling functions for each of the five mission areas. The *Global Channels Directorate* is responsible for all channel route missions. The *Mobility Management Directorate*, referred to as the *barrel*, oversees the allocation of aircraft to the different mission areas by continually balancing the mission area priority with the available number of aircraft and aircrews. During mission execution the *Command and Control Directorate*, referred to as the *floor*, tracks each mission minute-by-minute and provides real-

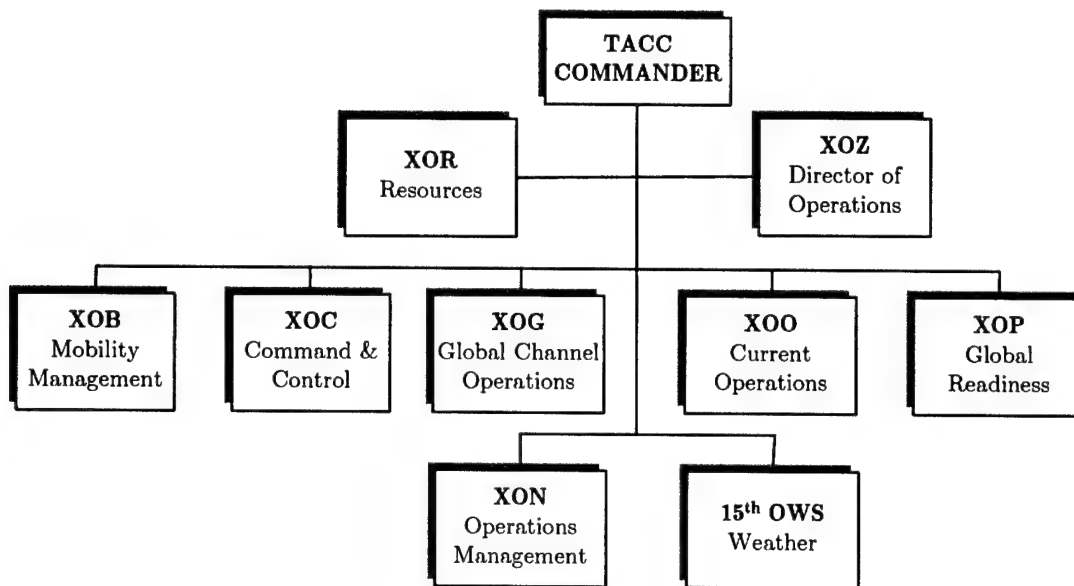


Figure 2-2: TACC organizational chart

time re-planning when necessary. The remaining directorates handle all the information required to plan and execute an aircraft mission, ranging from weather to international diplomatic and over-fly clearances [50].

The Global Channels Directorate is divided into an East Division and West Division. The West Division is responsible for all the channel route missions that originate at locations on the western coast of the U.S. and at Charleston Air Force Base (AFB), South Carolina. Although Charleston AFB is located on the east coast, it falls under the West Division so that the workload between the two approximately equally manned divisions is balanced. The East Division is responsible for all remaining channel route missions that originate on the eastern coast of the U.S.

In its entirety, the TACC is the operational arm of AMC, acting as the sole organization that schedules and overseas every mission flown by strategic airlifters in the DTS.

2.1.3 AMC's Three Objectives

The purpose of this section is to introduce AMC's three operational objectives of *readiness*, *customer service*, and *NOR*. We discuss each objective in the context of how each influences channel routes.

2.1.3.1 Wartime readiness though peacetime preparation

At the outset of each fiscal year (FY) the United States Congress approves a *Flying Hour Program* (FHP) that must be implemented by AMC. The FHP specifies the number of hours each aircraft type must operate during the year. The required number of hours for each aircraft type is determined using two requirements. The first is to develop the skills and experience (gained through accumulating flying hours) of pilots for wartime. The second is to enable pilots' progression from co-pilots to aircraft commanders (based partly on the total number of hours flown) in order to keep the correct balance of the two as older pilots leave the cockpit and new pilots enter.

MISSION AREAS

AMC is responsible for five different *mission areas* of flying operations. The hours flown by aircraft in each of these mission areas count toward the FHP. Thus, AMC must determine how to allocate the annual number of flying hours set in the FHP among the five mission areas. The five mission areas are *contingencies*, *exercises*, *special assignment airlift missions* (SAAMs), *Joint Airborne/Air Transportability Training* (JA/ATTs), and *channel routes*.

Each mission area varies by the purpose it serves, the Joint Chiefs of Staff (JCS) Priority code it receives, and the amount of planning lead-time provided to the schedulers at AMC. A mission area's priority is important because the same pool of aircraft is used to conduct all of the mission areas. Thus, in the event that there are not enough aircraft to fly all missions in each mission area, the missions from lower priority mission areas will not be flown using military aircraft. Furthermore, the planning lead-times can be viewed as a surrogate of a mission area's predictability. A mission area's predictability influences how planners determine the number of the flying hours expected to be flown within each mission area. Mission areas that are predictable and have a high JCS Priority can be allocated a specific number of flying hours. However, mission areas that have a high JCS Priority and are unpredictable are difficult to assign a specified number of flying hours. Similarly, it is difficult to assign a specified number of flying hours to a mission area that has a low JCS priority and is predictable because planners cannot be assured that these mission areas will not be trumped by mission areas that have a higher JCS priority and are less predictable. For example, contingencies have a high JCS priority but are unpredictable and channel routes have a lower JCS priority but are predictable. Thus, relying on contingencies to provide a consistent percentage of total flying hours from year-to-year would cause AMC to underfly or overfly the FHP. This is illustrated in *Figure 2-3* by the significant decrease in the number of contingency flying hours from FY99 to FY00.




Figure 2-3: Mission areas by percentage of total flying hours for FY99 and FY00

Contingencies

Examples of military contingency operations include the United States' recent involvement in Bosnia and Afghanistan. All air mobility missions flown in support of contingency operations are classified as *contingency* missions. By their very nature, *contingency* operations are highly unpredictable but have an extremely high priority. Consequently, unless a contingency is ongoing at the time the FHP is approved, planners do not have the ability to accurately assign a specified number of flying hours to be fulfilled through contingency missions. As a rough estimate, planners use conservative estimates of the number of flying hours that have historically been flown by contingency missions.

Exercises

The different armed services (i.e., USA, USN, USAF, U.S. Marine Corps) frequently gather in a common location to train with one another as well as with the military forces from other nations. These events take months, and even years, of planning and coordination in order to bring together all the necessary personnel and equipment. AMC is responsible for moving all the required U.S. military personnel and equipment for these events. All the airlift missions in support of these activities fall into the *exercises* mission area. As a result of the required planning horizon for these events, exercises are a predictable source of flying hours. Furthermore, exercises have a high priority.

SAAMs

SAAMs are missions in which a customer pays to rent an entire aircraft to perform special service between locations that are not included in established AMC routes. USTRANSCOM assigns these missions based on the number of passengers, the characteristics of the cargo (e.g., weight, size), and the urgency of the requirement [51]. These missions are highly unpredictable and can be requested with little lead-time. SAAMs have a high priority. Similar to contingencies, the number of flying hours expected to be flown by SAAMs is based on estimates calculated from historical data.

JA/ATTs

JA/ATT operations are similar to exercises. The key difference is that the focus of JA/ATT missions is to practice a specific war fighting skill that requires two or more U.S. military services. Exercises, on the other hand, simulate an entire set of skills over longer time periods (ranging from several days to several weeks). JA/ATT missions most frequently occur between the Air Force and the Army. For instance, Air Force aircraft will airdrop the Army's 101st Airborne Division. Just like exercises, JA/ATT's take prior planning and coordination, which allows AMC planners to allocate a specified number of flying hours to the operations in this mission area. Furthermore, JA/ATT's have a lower priority than contingencies, but a higher priority than channel routes.

Channel Routes

A *channel route* is recurring service between an origin and a destination. Channel routes are used to move cargo and passengers throughout the world and can be viewed as peacetime *distribution* and *sustainment* missions. These missions ensure that military personnel throughout the world, which have readiness requirements similar to pilots, have the resources required to maintain peacetime readiness. Channel routes are meant to provide predictable service, but they commonly have the lowest priority among the mission areas. Consequently, it is difficult for planners to know how many flying hours to allocate to channel routes. *Table 2-1* summarizes the priority, the predictability, and the percent of total flying hours by mission area.

For the remainder of the thesis, the following terminology will be used. A *channel* refers to a validated origin-destination pair (the validation process is discussed later in this section). *Channel route* refers to the mission area. The *channel route network* refers to the aircraft, bases, and cargo that are associated with any *channel route mission* in the *channel route schedule*. The *channel route schedule* is a monthly schedule of all *channel route missions*. A

	JCS PRIORITY (1 = highest)	PREDICTABILITY	% OF TOTAL FY99 FLYING HOURS	% OF TOTAL FY00 FLYING HOURS
<i>Contingencies</i>	1	Low	27%	11%
<i>Exercises</i>	2	High	6%	6%
<i>SAAMs</i>	3	Low to Medium	16%	23%
<i>Channel Routes</i>	4	High	27%	31%
<i>JA/ATTs</i>	5	High	25%	29%

Table 2-1: Priority, predictability, and percent of total flying hours by mission area


channel route mission (also referred to as a mission) specifies the *route*, specific timing of the route, and the flying squadron (discussed in §2.2.3.1) that will fly the channel route mission. A *route* is a sequence of stops, crew configuration, and aircraft type to be flown. *Channel route planners* consist of all the personnel that create the *channel route schedule* and the *channel route planning process* refers to all the activities required to generate a *channel route schedule*.

There are two types of channel routes: *requirements channels* and *frequency channels*. Requirements channels are scheduled solely on the actual volume of cargo located at a channel's origin, and frequency channels are scheduled solely on the frequency specified in the validated requirement. A frequency channel is scheduled regardless of the amount of actual cargo to be moved whereas a requirements channel will not be scheduled until enough cargo exists to justify an aircraft mission. According to USTRANSCOM, "frequency channels are established when traffic volume does not support the desired frequency of service. These channels support the operational necessity and quality of life requirements in remote areas" [51]. Frequency channels have a higher priority than requirements channels.

All channels must be *validated* by USTRANSCOM. The validation process is initiated when a customer submits a request for transportation to USTRANSCOM. Included in the request is the type of cargo to be moved, the origin and destination of the cargo, and the amount of cargo and/or the frequency of service desired (i.e., an aircraft with 25 tons of capacity once a week). From this information, USTRANSCOM *validates* the request by determining the appropriate mode of transportation (air, land, sea, or a combination of the three) and then sends the validated requirement to the organization responsible for scheduling the assigned mode of transportation. For frequency channels, USTRANSCOM also determines the *best-fit aircraft* for the request based on the type of cargo and the customer's estimate of the amount of cargo. The best-fit determination defines the *underutilization rate* that the customer is subject to (discussed in §2.1.3.3). If the assigned mode of transportation is air, USTRANSCOM will post the validated channel to the *Sequence Listing* and pass it to AMC.

The *Sequence Listing* is a list of all the validated requirements channels and frequency channels. It lists the origin and destination of the channel, the customer, the type of channel (i.e., requirements channel or frequency channel), and the frequency of service requirement if the channel is a frequency channel. The Sequence Listing is published annually and revised quarterly.

The channels in the Sequence Listing are not merely serviced by direct, point-to-point flights. Rather, channel route planners attempt to create aircraft routes (i.e., sequences of stops) that service several channels. The Sequence Listing serves as a list of allowable stops for all aircraft routes and of frequencies of service that must be met. At each stop in an aircraft route there must be a validated channel between the current stop and a stop prior and/or after the current stop. *Example 2-1* illustrates this concept.



Example 2-1: Consider four stops, A, B, C, and D. Assume the Sequence Listing contains the requirements channels (A,B) and (A,C) and the frequency channel (B,C) with a 1/week frequency of service. The route $A \rightarrow B \rightarrow C \rightarrow A$ is an allowable route because each stop is the origin (destination) of a validated frequency and/or requirements channel whose destination (origin) is a future (prior) stop in the route. An aircraft route $A \rightarrow B \rightarrow D \rightarrow A$ would not be permitted because there does not exist a requirements and/or frequency channel either between stop A and stop D or between stop B and stop D.

2.1.3.2 Recovery of operating costs

There are two funds that pay for AMC's flying operations – the *Transportation Working Capital Fund* (TWCF) and the *Aircraft Readiness Account* (ARA). The TWCF is a “revolving industrial fund for defense transportation” and is “USTRANSCOM's primary source of operating and capital investment funding” [52]. The TWCF is funded through the rates that customers are charged for transportation services. The ARA is AMC's readiness account (i.e.,

training budget). The ARA is used to subsidize customer rates and is funded through direct appropriation from the U.S. Congress. The size of the ARA is the difference between the revenue required to recover all operating costs and the revenue generated from customer rates [51].

The TWCF – which is managed by USTRANSCOM – is designed to recover all aircraft operating costs. In order to help recover all operating costs, AMC subsidizes the rates charged to customers for the different mission areas. The rates vary depending on the mission area. Channel route rates vary depending on the type of cargo shipped and are designed to be competitive with commercial rates for equivalent cargo. AMC subsidizes the amount required to recover full operating costs. The justification of this subsidized payment structure is that aircrew training is a byproduct of moving cargo for the customer. Thus, the customer only pays to move the cargo, and AMC pays for the training (from the ARA) that takes place as a result of transporting the cargo. Contingencies are unique. These costs are calculated independently so that peacetime customers do not bear the responsibility of funding large mobilizations. Contingency costs are funded through direct appropriation from the U.S. Congress, rather than through reimbursement of the TWCF [51].

The *TWCF customer rates* are determined eighteen months prior to the FY they are implemented and remain fixed for the entire FY. USTRANSCOM calculates TWCF rates based on expected workload (i.e., the amount of cargo and passengers expected to be moved), expected costs (e.g., fuel), and any prior gains or losses in the *Net Operating Result* (discussed in detail later in this section). After the rates have been determined, the *TWCF Planning Weights* are calculated. These weights, which are specific to each aircraft type, are the tonnages that must be moved for full recovery of the operating costs given the revenue income from the customer rates and ARA subsidies. For instance, assume that a C-5 aircraft must be loaded with fifty tons of cargo to recover the operating costs. Channel route planners will attempt to put at least fifty tons of cargo on the C-5 when they build the channel route network. If the C-5 does not have at least fifty tons of cargo when the mission is executed, the difference must be reconciled using the ARA. In other words, the TWCF Planning Weights for each aircraft serve as a lower limit on the amount of cargo that channel route planners attempt to load on the aircraft.

Frequency channel customers pay a higher rate than requirements channel customers. Frequency channel customers are also subject to an *underutilization rate* because the frequency channel route mission will be flown regardless of the amount of cargo actually moved. The *underutilization rate* determines the amount a customer must pay if the amount of cargo shipped for that customer is less than the estimated cargo weight that the customer submitted

to USTRANSCOM for validation. This rate is based on the best-fit aircraft determined by USTRANSCOM in the validation process. If the channel route planners assign an aircraft other than the best-fit aircraft to fulfill the frequency channel (possibly because requirements channels are also being satisfied with the channel route mission), the customer will be charged the underutilization rate associated with the best-fit aircraft determined during the validation process. If the amount of cargo shipped is greater than the amount the customer estimated in the validation request, the subsidy from the ARA is reduced. Furthermore, in the case that a channel route mission covers multiple frequency channels and requirements channels, the customer is only charged the underutilization fee for the flight leg that ends at the frequency requirement's destination, not over the entire channel route mission.

The *Net Operating Result* (NOR) is defined to be the revenue required to break even (i.e., recover all aircraft operating costs) less the ARA subsidies and the revenue generated by the TWCF customer rates. AMC's goal is to have a NOR of zero. Because TWCF customer rates are fixed eighteen months prior to the time period they are actually used and based on *estimated* workload and costs, NOR can be positive or negative. NOR will be positive (i.e., a profit will be realized) when the actual workload is greater than expected and/or estimated costs are greater than actual costs. The profit is then used to lower the TWCF customer rates for the following year. However, when workload is less than expected and/or estimated costs are less than actual costs, NOR will be negative. Again, this will be reflected by an increase in the following year's TWCF customer rates. In theory, the ARA can be reduced when an aircraft carries more than its TWCF planning weight. In order to reduce the ARA, AMC must achieve a worldwide utilization rate (annual utilization rate of all aircraft missions) of approximately 48%. The FY00 worldwide utilization rate was near 42% [43].

The FHP also influences NOR. Lower than expected workload results in less revenue producing cargo movements. The FHP requirements remain constant, however, regardless of the amount of cargo that must be moved. As a result, it is not uncommon for channel route missions to fly empty – without moving any revenue producing cargo – to meet the FHP requirements. A large number of non-revenue channel route missions result in a negative NOR. AMC does not exist to make a profit. To keep NOR near zero, however, it is beneficial for AMC to plan the channel route schedule to recover as much of the operating costs as possible, especially in years when workload is less than expected.

2.1.3.3 Customer Service

The number of contingency operations (e.g., the Gulf War, Bosnia) increased significantly in the 1990's compared to earlier decades. Simultaneously, the U.S. military was

transitioning to a post Cold War posture. A consequence of this transition was a significant decrease in the number of overseas military installations. Both factors had a significant impact on the FHP and channel routes.

As a result of the high number of contingency operations during the 1990's, contingency missions were used to make up any shortfalls in flying hours that could not be met with the other mission areas. Furthermore, the large number of contingency missions consumed nearly all the available military aircraft. Consequently, lower priority missions such as channel route missions, were not flown using military aircraft. This forced AMC customers to rely on commercial airlift to deliver the cargo normally moved by military airlift via channel routes. These commercial carriers were able to provide reliable service to customers served via channel routes. Now that the number of contingencies has decreased, AMC would like channel routes – flown by military aircraft – to become the backbone of the FHP. Channel routes are an ideal candidate because they are predictable and thus have the potential to be a stable source of flying hours. However, because channel routes have a low priority, it is a challenge to maintain the level of service that customers grew accustomed to with commercial carriers.

Another challenge facing channel route planners is the reduction of overseas military installations. Fewer overseas military installation results in less cargo to be moved throughout the world. The level of peacetime readiness (i.e., the number of required flying hours) has remained relatively constant, however. This was not a problem when numerous contingency operations could be used to satisfy flying hours requirements, but less cargo coupled with fewer contingencies makes the FHP a challenge for AMC.

Although AMC's top priority is to ensure the readiness of pilots, AMC's customers must also meet their readiness requirements. An integral piece of AMC's customers' readiness is timely delivery of personnel, equipment, and supplies. Thus, it is imperative that AMC provide predictable, reliable service. AMC must employ careful planning in order to maintain its own readiness while providing the customer service required for other military organizations to maintain readiness.

2.2 The Physical System

The following sections describe the physical entities in the channel route network, which include the *cargo items* and *aircraft* that flow through the system over time and the static *infrastructure* located at the *aerial ports* (military equivalent of airports).




Figure 2-4: Channel cargo shipped by type for FY99 and FY00

2.2.1 Cargo Items

There are a total of twenty-seven different types of cargo that flow through the channel route network, ranging from aircraft parts to rations and subsistence. In FY00, AMC shipped 163,174 short tons (one short ton equals 2,000 pounds) of cargo, down from the FY99 total of 219,994 short tons. The nine cargo types that accounted for most of the cargo shipped via channel route missions are presented in *Figure 2-4*.

2.2.1.1 Cargo Item Categories

The cargo moved via channel routes is characterized in one of three general categories: *bulk*, *oversized*, or *outsized*. Bulk cargo fits on a standard 463L cargo pallet, which can be unloaded to an aircraft and offloaded from an aircraft using a cargo loader (see *Figure 2-5*). The length and width of a standard pallet are 108 inches and 88 inches, respectively. The height of a pallet varies depending on the aircraft load configuration and the nature of the cargo, but the maximum height of a pallet is 96 inches.




Figure 2-5: Standard pallet and cargo loaders

A pallet built to standard dimensions fits on any aircraft type. As a result of the varying heights of pallets, the weight of a pallet can vary greatly. Although a standard pallet will physically fit on each aircraft type, the weight and balance limitations of each aircraft type differ and can influence the ability of a pallet to be transferred between aircraft.

Cargo that cannot fit on a single standard pallet is categorized as oversized cargo. In many cases, two or more standard pallets are combined to ship oversize cargo. Outsize cargo is large, unwieldy cargo such as tanks and helicopters. Bulk cargo accounts for approximately 99% of all cargo moved via channel routes. This greatly simplifies matters for channel route planners because all cargo with a common origin aerial port and destination aerial port can be aggregated to determine the total number of short tons that must be moved between the origin and the destination. This alleviates the need for channel route planners to decompose the cargo into the three cargo categories.

Each item of cargo is assigned a priority, which determines the order in which it is moved from the *aerial port* (discussed in §2.2.2). The governing cargo-processing rule at the aerial ports is “first-in-first-out (FIFO) by priority”. The cargo priorities, from highest to lowest, are 999 (“triple 9”), TP-1, TP-2, TP-3, and TP-4. The cargo priority is determined by USTRANSCOM after the customer submits a request for transportation. The cargo priority is one of the factors that determine the assigned mode of transportation. For example, TP-3 and TP-4 cargo will rarely be shipped via air transportation. Cargo items with these priorities will be shipped via truck and/or sealift.

2.2.1.2 Cargo Flow

Channel route cargo can follow various different paths as it flows through the channel route network, as illustrated in *Figure 2-6*. Cargo originating in the CONUS (*Figure 2-6a*) starts at one of many CONUS distribution warehouses located throughout the U.S.. The two

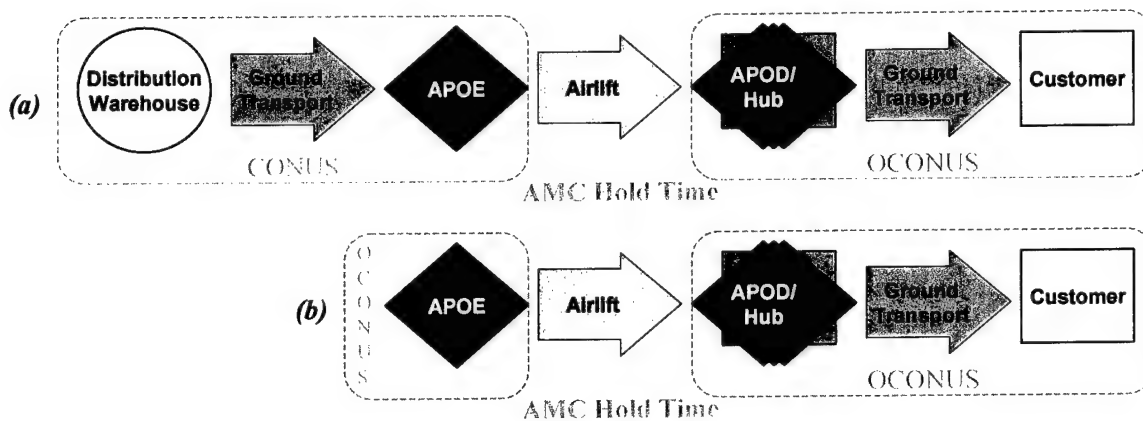


Figure 2-6: Cargo flows originating at (a) CONUS APOE, or (b) OCONUS APOE

largest *Defense Distribution Center* (DDC) facilities are located in Susquehanna, Pennsylvania and San Joaquin, California. There are a total of twenty-four DDC warehouses. Additionally, cargo can originate at one of the many armed services' supply depots, at one of eleven *Defense Commissary Agency* (DeCA) distribution centers, at commercial household goods movers' warehouses, or at one of the eight General Services Agency (GSA) distribution centers.

Cargo is trucked from the distribution warehouse to one of the CONUS aerial ports, called the *Aerial Port of Embarkation* (APOE). At the APOE, the cargo is palletized, if not already palletized at the distribution warehouse. Each cargo item is assigned a *Transportation Control Number* (TCN) at the APOE. The customer can use the TCN to track the location of the cargo item within the DTS. If the cargo originates in OCONUS (*Figure 2-6b*), it originates at an OCONUS APOE.

After assigning each piece of cargo a TCN and palletizing the cargo, the pallet is loaded onto an aircraft at the APOE and flown either directly to its final destination, called the *Aerial Port of Debarkation* (APOD), or through one or more enroute aerial ports. If the cargo is flown directly to the APOD it is offloaded and put either into a storage holding facility or loaded directly on a truck and delivered to the customer, if the customer is not located at the APOD. If the cargo is put into an APOD cargo holding facility it will eventually be loaded on a truck and delivered to the customer. If the cargo is not flown directly and one of the enroute stops is a hub (discussed in §2.2.2), the cargo will either be offloaded and put into a cargo holding facility, immediately onloaded to another aircraft and flown to the APOD, or it will remain on the aircraft. It should be noted that a hub can be an APOD. If the cargo is put into the hub's cargo holding facility, it will eventually be onloaded to another aircraft and flown to the APOD. Once at the APOD, the cargo is either trucked directly to the customer or placed in a cargo

holding facility and trucked at a later date, if the customer is not located at the APOD. If the enroute stop is not a hub, the cargo whose destination is not the enroute aerial port will remain on the aircraft. If there is cargo at the enroute aerial port that is destined for a future aerial port in the channel route mission, it will be unloaded to the aircraft.

The AMC hold time is the time period starting when the cargo is delivered to the APOE until the time it arrives at the APOD. AMC hold time is just one piece of *customer wait time*, which starts when a customer places an order and ends when the customer receives the shipment.

2.2.2 Aerial Ports

Aerial ports are the physical locations where cargo is consolidated from the distribution warehouses and where aircraft land, takeoff, onload/offload cargo, and refuel. Aerial ports contain runways, ramp space for aircraft parking and servicing, cargo holding facilities, cargo loading equipment, fuel storage, and aircraft fuel trucks. The infrastructure, servicing facilities, and manpower at each aerial port limit the number of aircraft the aerial port can handle. Some aerial ports also serve as the home base for one or more aircraft types (discussed in §2.2.3.1).

Aerial ports that serve as the origin (starting location) of a channel route are referred to as Aerial Ports of Embarkation (APOEs) and aerial ports that are the final destination are referred to as Aerial Ports of Debarkation (APODs). There are six CONUS APOE's and each traditionally services various regions throughout the world. *Figure 2-7* illustrates the locations of the CONUS APOE's and the regions associated with each. An arrow between an APOE and a region indicates that at least one validated requirements channel and/or frequency channel exist between the CONUS APOE and an aerial port in the region.

After the cargo has arrived at the APOE and has been assigned a TCN and palletized (if it is not palletized at the depot), it is placed into an indoor storage warehouse. These warehouses have extremely large capacity and rarely ever approach maximum capacity during peacetime. Consequently, channel route planners do not consider APOE warehouse capacity when creating a channel route schedule.

Aerial ports also house all the necessary equipment, called *servicing resources*, required to build pallets, transport pallets to and from the aircraft, and refuel the aircraft. The number of servicing resources available at an aerial port determines the number of aircraft that can simultaneously be serviced, called *Working Maximum on Ground* (WMOG). This is an important planning consideration for the channel route planners. Furthermore, the total number of aircraft at an aerial port cannot exceed the *Parking Maximum on Ground* (PMOG)

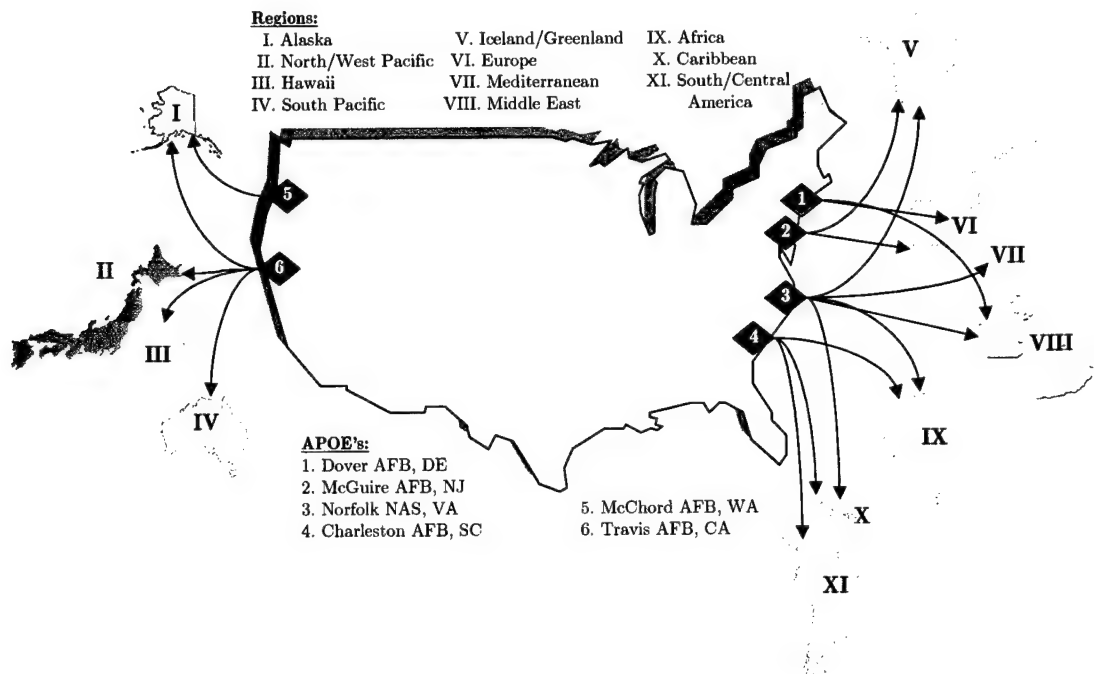


Figure 2-7: CONUS APOEs and associated regions throughout the world

capacity. PMOG is a reflection of the physical ramp space (i.e., aircraft parking area) available at an aerial port.

In addition, each aerial port has associated *operating hours* (ops hours), *Bird Air Strike Hours* (BASH), and, in some instances, *quiet hours*. Operating hours limit the times aircraft can arrive and depart from an aerial port. These hours reflect when the aerial ports are manned. BASH hours reflect the times of day during each season of the year that aircraft are prohibited from landing or taking off at an aerial port because of heavy migratory bird activity. Quiet hours are the times of the day certain aircraft are prohibited from landing or taking off at an aerial port. Quiet hours are aerial port specific and do not exist at all aerial ports.

AMC has specified several OCONUS aerial ports as *hubs*. These hubs are modeled after the same hub-and-spoke concept that many major U.S. airlines and air cargo shippers use. At hub locations, cargo from various origins is sorted, consolidated, and delivered to common destinations. Each hub serves a specific geographic region. Currently, there are five OCONUS hubs in the channel route network, illustrated in *Figure 2-8*.




Figure 2-8: OCONUS hubs in the channel route network

2.2.3 Aircraft Assets

This section describes the military (organic) aircraft and commercial aircraft that are used to transport channel route cargo. The aircraft used to conduct channel route missions come from four sources: Air Force active duty, Air Force Reserves, Air National Guard, and commercial carriers that are *Civil Reserve Air Fleet* (CRAF) participants.

2.2.3.1 Organic Aircraft

There are several organic aircraft types, referred to as strategic airlifters, used to transport cargo: Lockheed Martin C-5 Galaxy, Boeing C-17 Globemaster III, Lockheed Martin C-141 Starlifter, Lockheed Martin C-130 Hercules. Each of these aircraft can carry bulk and oversized cargo. However, only the C-5 and the C-17 have the ability to move outsized cargo. Furthermore, the C-5, C-17, and C-141 have the ability to refuel while in flight, called aerial refueling (AR). AR significantly increases aircraft range capabilities, but introduces an additional level of complexity into the channel route planning process.

The Air Force active duty component operates all four types of strategic airlifters, whereas the Air Force Reserve component and the Air National Guard component operate all but the C-17. The Air Force Reserve and Air National Guard components are often grouped together and referred to as the *Air National Guard & Reserve Component* (ARC). In FY99, the ARC component was responsible for the operation of approximately fifty-eight percent of the total strategic airlift force [48].

The organic aircraft and aircrews are organized into *aircraft wings*. Each aircraft wing consists of one or more *aircraft squadrons*. An aircraft squadron consists of only one aircraft

type. Thirty days prior to the execution month, the aircraft wings will allocate a portion of the aircraft and aircrews in each aircraft squadron for TACC's use. These aircraft and aircrews are referred to as the *contract aircraft and aircrews* and provide an upper limit on the number of aircraft and aircrews that can be assigned by the TACC to fly channel route missions, contingencies, exercises, and SAAMs. The remainder of the aircraft in the aircraft squadrons will be used for local training missions. These aircraft are referred to as *fenced aircraft* and they cannot be assigned by the TACC. Although contract aircraft and aircrews are tasked and scheduled by the TACC, these resources must return the aircraft squadron's home base at the conclusion of a mission from any mission area.

Each aircraft has different capabilities, including cargo capacity and maximum flying range. Additionally, each aircraft has an associated *minimum ground time*. The minimum ground time is the time the aircraft is required to remain on the ground after landing, regardless of the amount of time it takes to perform offloading and onloading of cargo, refueling, and other servicing functions. The only situation that warrants a shorter ground time is when the aircraft lands solely to refuel. These characteristics are summarized in *Table 2-2*.

Each aircraft also has a different cruising speed. There are many factors other than an aircraft's cruising speed that determine the amount of time it takes an aircraft to fly between two points (e.g., weather and atmospheric winds). To deal with these factors during the channel route planning process, aircraft flight times are calculated using historical flight times. TACC planners use the *TACC Flight Time Calculator* to access a database that stores the actual flight times of all past AMC missions. The TACC Flight Time Calculator will calculate the average flight time for a specific aircraft type between two locations. The TACC Flight Time Calculator will also report the number of missions used to calculate the reported average flight time.

Aircraft <i>Manufacturer</i>	Max Takeoff Weight (tons of cargo & fuel)	TWCF Planning Weight (tons)	Min Ground Time (hrs+min) (Refuel Only)
C-5 Galaxy <i>Lockheed Martin</i>	145	50	4+15 (3+15)
C-17 Globemaster III <i>Boeing</i>	85	25	3+15 (2+15)
C-141 Starlifter <i>Lockheed Martin</i>	34	16	3+15 (2+15)
C-130 Hercules <i>Lockheed Martin</i>	22	8	2+15 (2+15)

Table 2-2: Organic aircraft characteristics

2.2.3.2 Commercial Augmentation

There are situations in wartime and peacetime when the military does not have enough airlift capability to move all the cargo that customers want moved. In these situations, commercial air carriers are contracted by AMC to make up this deficit. Participation in the wartime commercial augmentation program – the *Civil Reserve Air Fleet* (CRAF) – determines if a commercial carrier can bid for peacetime contracts through the *fixed-buy* and *expansion-buy* programs.

CIVIL RESERVE AIR FLEET (CRAF)

President Truman established the *Civil Reserve Air Fleet* (CRAF) in 1951 to augment military airlift capabilities in wartime and national mobilization efforts. In the CRAF program, commercial carriers pledge a certain percentage of their aircraft fleet and aircrews for military use when the President activates the CRAF program. In return, air carriers are paid a fixed contract price for the missions flown if the CRAF program is activated. The first time CRAF was activated was during the Gulf War, during which CRAF participants “flew 20% of all the airlift missions accounting for 70% of all passenger movements and 30% of all cargo shipments” [19].

Although commercial carriers do not receive money for the aircraft they commit to the CRAF program unless CRAF is activated, the number of aircraft that a carrier commits also earns that carrier a percentage of the available peacetime bid points. These bid points can then be used to procure two types of peacetime contracts, *fixed-buy contracts* and *expansion-buy contracts*, which are offered exclusively to CRAF participants. The purpose of these two programs is to augment military airlift in peacetime. Specifically, each carrier can bid on any fixed-buy and expansion-buy contracts they have the resources to execute. The fixed-buy and expansion-buy contracts are awarded to the commercial carriers based solely on the number of peacetime bid points. Thus, more peacetime bid points results in more peacetime contracts. This entices the commercial carriers to commit more aircraft to CRAF by awarding them the opportunity for more peacetime contracts. *Table 2-3* lists the FY00 CRAF participants.

COMMERCIAL FIXED BUY PROGRAM


At the outset of each year, AMC determines the amount of peacetime commercial augmentation required to meet the peacetime airlift requirements. To determine the amount of commercial airlift that is expected, AMC will look at historical fixed-buy contracts, historical cargo data, and the FHP. AMC is careful not to contract too much commercial airlift because

Air Transport International	Express One International	Reeve Aleutian Airways
Alaska Airlines	Federal Express	Southwest Airlines
American Airlines	Fine Air	Spirit Airlines
American Trans Air	Gemini Air Cargo	Sun Country Airlines
America West	Kitty Hawk International	Sunworld Int'l Airlines
Arrow Air	Lynden Air Cargo	Trans World Airlines
BAX Global	Miami Air	Tower Air
Continental Airlines	Midwest Express	United Airlines
Delta Airlines	North American Airlines	United Parcel Service
DHL	Northern Air Cargo	U.S. Airways
Emery Worldwide	Northwest Airlines	U.S. Airways Shuttle
Evergreen International	Omni Air International	World Airways
	Polar Air Cargo	

Table 2-3: FY00 Civil Reserve Air Fleet (CRAF) Participants

there must be enough flying hours available for military aircraft to meet to the FHP requirements. However, at the same time, AMC would like to guarantee its customers a predictable level of service and give the commercial carriers enough business to make participation in CRAF worthwhile. Military aircraft cannot always provide a guaranteed level of customer service via channel route missions because of the higher priority mission areas.

Determining the routes to offer as fixed-buy contracts is not trivial. AMC recovers more of the operating costs with channel route missions that have a higher utilization, thus reducing the amount AMC must reconcile using the ARA and, ultimately, helping to prevent a negative NOR. However, commercial carriers are paid on a per ton-mile basis. Thus, commercial carriers must have a relatively large amount of cargo to move in order to profit. *Example 2-2* illustrates this concept.



Example 2-2: Consider two channel route missions: $A \rightarrow B \rightarrow C \rightarrow A$ and $A \rightarrow D \rightarrow E \rightarrow A$. Let b_{ij} represent the amount of cargo at location i that is destined for location j . Let $b_{AB} = 25$, $b_{AC} = 25$, $b_{BC} = 25$, $b_{AD} = 5$, $b_{AE} = 10$, & $b_{DE} = 5$. Assume that the organic aircraft operating these routes has a TWCF planning weight of 50 tons,

which recovers 70% of the operating costs. Assuming an aircraft carries 50 tons, this leaves AMC to subsidize 30% of the operating costs through the ARA. Any unused capacity below 50 tons will result in a subsidy (funded from the ARA) greater than 30%. For the $A \rightarrow B \rightarrow C \rightarrow A$ route, the aircraft will carry 50 tons of cargo from $A \rightarrow B$. At B the aircraft will offload the 25 (A,B) tons and onload the 25 (B,C) tons. The amount carried from $C \rightarrow A$ will be 50 tons, and the ARA will only be required to subsidize the expected 30%. For the $A \rightarrow D \rightarrow E \rightarrow A$ route, the aircraft will carry 15 tons of cargo from $A \rightarrow D$, leaving 35 tons of unused capacity. At D, the aircraft will offload the 5 (A,D) tons and onload the 5 (B,C) tons. 10 tons will be carried from $D \rightarrow E$, leaving 40 tons of unused capacity. Consequently, AMC must pay much more than 30% of the operating costs for this route. Thus, it is beneficial, in terms of NOR, for AMC to use military aircraft over the routes with a large amount of tonnage. However, because the fixed-buy contract is based on a per ton-mile rate, the $A \rightarrow D \rightarrow E \rightarrow A$ route is less likely than the $A \rightarrow B \rightarrow C \rightarrow A$ route to entice commercial carriers. Thus, routes appealing to AMC are also the routes that are appealing to CRAF carriers.

Unfortunately, there are no analytic tools that AMC can use to evaluate the influence of different commercial fixed-buy structures on NOR and the FHP. An advantage of the optimization models developed in later chapters is that they can be solved quickly, which would allow various fixed-buy structures to be analyzed.

After AMC finalizes the amount of augmented commercial airlift they expect will be required, they will enter contract negotiations with the commercial carriers to determine the *fixed-buy rate*. The fixed-buy rate is the rate the commercial carriers will be paid per ton-mile for the cargo they move under the fixed-buy contracts they are awarded. To award the routes, AMC posts the requirements for the route, including the specific sequence of flight legs to be flown, the required capacity over the route, and the frequency the route must be flown. As an example, a possible requirement could be for 60 tons of capacity every Saturday and Tuesday from Dover AFB, Delaware to Ramstein AB, Germany to Mildenhall AB, England and back to McGuire AFB, New Jersey. Note that the contract does not specify an aircraft type, only the capabilities required. The commercial carriers determine the specific fixed-buy contracts that interest them. If there are multiple carriers that express interest in the same fixed-buy contract, AMC awards the route based on the number of peacetime bid points possessed by each commercial carrier. For instance, suppose two commercial carriers express interest in the contract mentioned above. If both carriers have roughly the same percentage of CRAF points, each carrier will be awarded one of the days, either Tuesday or Saturday.

Each fixed-buy contract is awarded for the duration of the fiscal year. Thus, using the fixed-buy contract example from the previous paragraph, the channel route planners are guaranteed that every Tuesday and Saturday a commercial carrier will provide 60 tons of capacity from Dover to Ramstein to Mildenhall and back to McGuire. This is important because unlike organic airlift resources, these aircraft cannot be tasked with higher priority missions. It is also important to note that the commercial carriers are responsible for providing the required aircrews for each contracted aircraft and for all maintenance issues. In fact, part of the contract specifies a reliability rate that the commercial carriers must meet. If a commercial carrier's reliability rate falls outside these limits a fine is imposed.

In FY01, the fixed-buy program included seventeen routes that were flown by four wide-body aircraft (e.g., Boeing 747), three narrow-body aircraft (e.g. L100), and ten combination aircraft (i.e., DC-8 aircraft configured to carry cargo pallets and passengers). The value of the entire contract was approximately \$163 million [41].

COMMERCIAL EXPANSION BUY PROGRAM

In addition to negotiating the fixed-buy rate, the commercial carriers and AMC also negotiate the *expansion-buy rate*, which is greater than the fixed-buy rate. This is the rate the military will pay the commercial carriers for short notice contracts, called *expansion-buy contracts*. The need for these contracts arises when there are not enough aircraft to move all the required cargo in the channel route network. This can happen for a number of reasons, including unexpected spikes in cargo generation at a specific aerial port or the loss of an organic aircraft to a higher priority mission.

Similar to the fixed-buy contracts, expansion-buy contracts specify the amount of cargo to be moved, the specific route to be followed, and the delivery date. The contract will not have a dollar amount specified. Rather, the commercial carrier will be paid accordingly to the expansion-buy rate. Historically, expansion-buy contracts will cost AMC an estimated five to ten million dollars each month [41].

2.2.4 Aircrews

Aircrews, like aircraft, have operating restrictions. Specifically, the aircrew must not exceed the *crew duty day* (CDD) limit, which is the number of consecutive hours a crew can fly. The CDD starts when a crew arrives at the APOE to plan the mission and inspect the aircraft, and ends when the aircrew arrives at the APOD. Between each CDD an aircrew must be given a minimum period of *crew rest*. There are two types of crews, *basic aircrews* and *augmented aircrews*. An augmented crew is a combination of two basic crews, which allows the aircrew to

Aircraft <i>Manufacturer</i>	Basic CDD (hrs+min) <i>(Augmented CDD)</i>	Min Crew Rest (hrs+min)
C-5 Galaxy <i>Lockheed Martin</i>	16+00 <i>(24+00)</i>	17+00
C-17 Globemaster III <i>Boeing</i>	16+00 <i>(24+00)</i>	16+30
C-141 Starlifter <i>Lockheed Martin</i>	16+00 <i>(24+00)</i>	16+00
C-130 Hercules <i>Lockheed Martin</i>	16+00 <i>(18+00)</i>	16+15

Table 2-4: Aircrew characteristics

extend its CDD. The aircrew characteristics are summarized in *Table 2-4*. Similar to aircraft, the flying wings will inform the TACC thirty days prior to the execution month of the number of aircrews that will be available.

It is typical for an aircrew to remain with an aircraft for an entire channel route mission, which typically spans anywhere from one to five days. There are a few select locations that aircrews have been prepositioned in order to quick-turn an aircraft. In these situations, an aircrew at the end of its CDD will begin crew rest at a selected aerial port (usually a hub) and a crew at the beginning of its CDD will assume flying the aircraft for the remainder of the channel route mission. When the original crew has fulfilled its crew rest requirements, it will assume the flying responsibilities of the next aircraft scheduled for a quick-turn at that aerial port. Prepositioned aircrews reduce the amount of time an aircraft sits idle; however, prepositioning crews adds another layer of complexity to the channel route planning process.

2.3 The Current Planning & Execution Process

Historically, AMC has not used scheduling algorithms or analytic models focused on peacetime scheduling on a day-to-day level. Models developed have either been for strategic analysis of large-scale deployments of personnel and equipment (e.g., Operation Desert Storm) or for a specific function such as aircraft load planning or aircraft assignment. While some of these models take channel route activity into consideration, channel route planners do not use these models for daily planning. There are several support tools that channel route planners use to facilitate the channel route planning process; however, none of the tools currently incorporate any type of scheduling algorithm. As a result, the channel route planning process is a largely

manual, iterative process. The channel route planners consist of the *organic channel schedulers*, *commercial schedulers*, *cargo bookies*, and *barrelmasters*.

This section reviews several of the analytic models that have been developed for use at AMC. This section also gives an overview of the TACC and details the functions of those involved in creating and overseeing the execution of the channel route schedule.

2.3.1 Literature Review – Historical Models at AMC

In 1985 the Air Force Logistics Management Center began an initiative to automate many of the manpower intensive planning and execution tasks associated with airlift operations. One of the earliest efforts of this initiative was the Deployable Mobility Execution System (DMES). DMES uses a cutting stock formulation to determine how to efficiently place cargo on an aircraft subject to volume, weight, and balance limitations of the aircraft [21]. This model requires predetermined cargo flows as an input.

In 1986 MAC's (AMC's predecessor) Director of Operations created an analysis group consisting of twenty-four operations research analysts to "provide analytical support to the Commander and his senior staff" [35]. During this time, the main analytic tool was the Mobility Analysis Support System (MASS). The core element of MASS is the Airlift Flow Model (AFM), which is a discrete event simulation model written in FORTRAN. This model has evolved and continues to be a favorite of analysts at AMC. Another tool used to analyze MAC's distribution systems during this time was an AT&T KORBX computer running the Karmikar algorithm [35]. Beginning in 1996, the Naval Post Graduate School (NPS) and the RAND Corporation combined several smaller models to create the NPS/RAND Mobility Optimizer (NRMO). NRMO is a large-scale, multi-commodity network based linear program that models the same system as MASS. Both are intended to model strategic airlift in the context of large military deployment scenarios such as Desert Storm [5],[54]. Both models consider channel routes as an input and do not attempt to design the peacetime channel route network.

In 1987, the development of the Airlift Deployment Analysis System (ADANS) began. ADANS is an interactive database decision support system that is used to schedule and analyze airlift missions as well as distribute the schedule to the flying units throughout the world [18],[31]. Work started on several peacetime planning modules in the late 1980's, but the wartime scheduling piece of ADANS had not been developed prior to the Gulf War. In 1990, the wartime portion of ADANS became the priority and was accelerated in order to aid planners with scheduling airlift operations for the Gulf War.

Included in ADANS is a two-step heuristic for generating airlift missions, called the Airlift Planning Algorithm (APA). The first step of the APA assigns cargo and passenger movement requirements to an aircraft by minimizing the unused aircraft capacity [30]. The second step of the APA schedules routes (i.e., determines the sequences of stops) such that the amount of time between pickup and delivery of the cargo load is minimized and all operational rules of the aircraft, aircrews, and aerial ports are satisfied [31].

Another tool within ADANS is an algorithm based on an insertion heuristic that is used to update an existing airlift schedule on a daily basis in order to incorporate new cargo requirements or changes in priority to the existing cargo requirements. This sequential heuristic attempts to minimize the number of changes to an already established schedule by adding cargo requirements one at a time to the schedule [47].

A third tool within ADANS uses a transportation problem formulation as a preprocessor to the APA. The formulation serves to prioritize the input list of cargo and passenger requirements, which is then input to the APA. Additionally, this formulation is used to estimate the total amount of cargo and number of passengers that can be moved using the APA [32].

ADANS does not attempt to design the peacetime channel route network. ADANS does, however, contain a linear programming formulation that determines the optimal flow of cargo through an existing channel route network. This formulation relies on a set of channel route missions as an input [40]. It is not well documented if and how the scheduling algorithms with ADANS are currently being used by the TACC to schedule daily aircraft operations.

The Strategic Transport Optimal Routing Model (STORM) is a linear programming model developed by strategic planners at AMC to determine the optimal cargo flow over a given set of input channel routes (i.e., sequences of flight legs). STORM determines the flow of cargo and aircraft that minimizes the costs of cargo handling, military aircraft operations, and commercial aircraft leasing. The solution provided by STORM is non-integer and is too large to solve efficiently using branch and bound. As a result, the variables are rounded down to the nearest integer, which results in unmoved cargo and unsatisfied frequency requirements. A heuristic based on the Effective Gradient Algorithm developed by Senju and Toyoda [1] is used to add single aircraft routes until all the unmoved cargo is shipped and the unsatisfied frequency requirements are satisfied. A major drawback of this formulation is that it lacks the ability to model aircraft movements over time and instead models the problem by aggregating an entire month's data. To assign timing to the routes, a simulation is used to schedule the chosen routes determined by the STORM model.

Another tool that has recently been tested at AMC is the *Barrel Allocator*. This scheduling tool is used for the day-to-day allocation of aircraft to the five mission areas. The problem is formulated as a fleet assignment problem. An initial solution is generated and then aircraft are reallocated to incorporate higher priority missions as they arise [14]. This model uses the schedules from each of the mission areas as an input. If all the missions cannot be flown the lower priority missions are changed (or dropped) and feedback is provided to the appropriate mission area planners. This formulation does not attempt to determine either the routing of a channel route mission or the cargo that should be carried by a channel route mission. The Barrel Allocator only determines if an aircraft scheduled to fly a channel route mission is available to actually fly the mission, considering the schedules from other mission areas.

2.3.2 Channel Route Planning Tools

The Global Decision Support System (GDSS) is the overarching command and control system for all airlift and aerial refueling activities. GDSS encompasses several databases and planning tools used by the TACC as well as by the aircraft squadrons. With respect to channel routes, GDSS serves as a database to house all planned channel route missions, including all execution details for each channel route mission. The execution details for a channel route mission include the specific route to be flown, takeoff and ground times at each location, the type of aircraft that will fly the mission, the flying wing that the aircraft and aircrew will come from, event schedule (i.e., onload, offload, crew rest, refuel, crew change) at each aerial port along the route, and any other mission specific remarks such as required diplomatic clearances and overfly restrictions.

Channel route planners save (and update, when necessary) all the channel route missions in the GDSS database. Aircraft squadrons throughout the world have access to the GDSS database, which allows them to determine the channel route missions they have been assigned to fly.

The Consolidated Air Mobility Planning System (CAMPS) is the primary planning tool that channel route planners use on a daily basis to create and update channel route missions. Furthermore, CAMPS serves as the scheduler's interface with GDSS. Schedulers have the ability to create and modify channel route missions directly within GDSS; however, the CAMPS graphical user interface makes it much more user friendly.

CAMPS calculates and verifies all appropriate ground times, crew rest periods, and crew duty day limitations. It does not currently calculate accurate flight times because it uses direct flight distances. Direct flight distances do not capture factors such as atmospheric winds and

international overfly restrictions. To overcome this inaccuracy, schedulers use the TACC Flight Time Calculator.

CAMPS significantly reduces the laborious tasks a scheduler faces of trying to gather and consolidate all the appropriate and critical information required to create a channel route mission. However, CAMPS does not currently incorporate automated scheduling algorithms that select and/or create the channel routes and determine cargo flows. As a result, determining the sequence of stops that comprise a route and the type of aircraft assigned to fly the route is a manual process. Similarly, when channel route missions are drastically changed because of aircraft and/or aircrew availability, CAMPS does not provide any decision support to select alternate channel route missions or to provide insights into how the changed channel route mission will influence the entire channel route network.

2.3.3 The Planning Process

The channel route planners within the TACC are responsible for creating and maintaining a monthly channel route schedule. The purpose of this section is to describe the current channel route planning process, which is largely a manual process that consumes many days. A general overview is presented followed by a detailed description.

2.3.3.1 General Overview

A channel route schedule consists of all the channel route missions scheduled for a calendar month. A different schedule is created for each calendar month. When that month arrives, it is referred to as the *execution month* (see Figure 2-9).




Figure 2-9: The channel route planning period and execution month

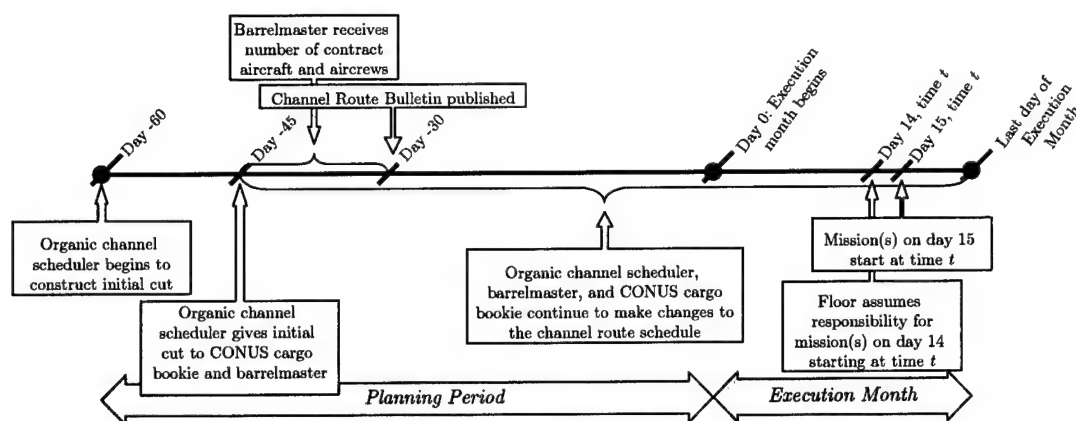


Figure 2-10: The channel route planning timeline of events

The planning period is initiated sixty days prior to the start of the execution month. During the planning period, channel route planners create all the channel route missions for the execution month and save them to the GDSS fifteen to thirty days prior to the execution month. Changes to the channel route schedule will continue through the remainder of the planning period and throughout the execution month. A timeline of scheduling activities discussed in the following sections and illustrated in *Figure 2-10*.

To begin the planning period, the *organic channel scheduler* accesses (via CAMPS) the channel route schedule that was created in the previous month and makes modifications to deal with any known special circumstances, such as planned runway construction and ARC component preferences. At this point, the schedule is referred to as the *initial cut*. The initial cut is then simultaneously passed to the *CONUS cargo bookie* and the *barrelmaster*. The initial cut is created without information about the *estimated cargo generation* and the contract number of aircraft and crews.

The barrelmaster looks at the schedules of other mission areas (e.g., SAAMs, exercises) and determines if there are enough available aircraft to fly the scheduled channel route missions in the initial cut. If there are not enough aircraft and aircrews to fly all the missions from all the mission areas, the barrelmaster will either *non-support* (i.e., drop) or significantly modify the lowest priority mission(s). Frequently, the lowest priority missions are the channel route missions that service requirements channels. The non-supported missions can be offered to the ARC component (who have no obligation to fulfill), offered as expansion-buy contracts to the commercial carriers through the commercial scheduler, or completely removed from the schedule. Non-supported channel route missions and significantly modified missions (e.g., entire

flight legs dropped or aircraft types changed) occur relatively frequently because channel routes must compete with mission areas that have a higher priority and short planning lead-times.

Using the same version of the initial cut passed to the barrelmaster, the *CONUS cargo bookie* determines if there are enough scheduled missions to move all the cargo expected to be in the system. If not, the CONUS cargo bookie will request that the organic channel scheduler create a new channel route mission. Before the channel route mission can be created, however, the barrelmaster must approve an aircraft to fly the new channel route mission. If an active duty organic aircraft is available, the new channel route mission is added to the channel route schedule. If, on the other hand, an active duty organic aircraft is not available, the new channel route mission will be offered to the ARC component, offered as an expansion-buy contract to the commercial carriers, or the cargo will not be moved.

To deal with new and non-supported channel route missions either the ARC component or commercial carriers can be used. The CONUS cargo bookie and/or the organic channel scheduler will contact the ARC component squadrons to determine if they are available to fly the mission. In order to offer a new or non-supported channel route mission as an expansion-buy contract, the cargo bookie must coordinate with the commercial scheduler and receive authorization from the TACC Commander and the AMC Director of Operations.

The current scheduling process is illustrated in *Figure 2-11*. *Figure 2-11* is provided as an overview and graphical representation of the details discussed in the following sections.

2.3.3.2 Organic Channel Scheduler

The organic channel scheduler is responsible for creating a list of channel route missions at the beginning of the planning period. The organic channel scheduler continues to maintain the schedule, making necessary changes to the channel route schedule throughout the planning period and the execution month. There are three organic channel schedulers to handle the scheduling responsibilities. At any given time, there are two organic channel schedulers in the planning period phase and one organic channel scheduler actively working with the execution months' schedule.

The planning process begins sixty days prior to the start of the execution month when the channel route scheduler begins to compile a set of aircraft missions using the initial cut from the previous month and, possibly (there is no standard method followed), the initial cut from the same calendar month from the previous year. Changes to this set of channel route missions are made to deal with known special circumstances. By forty-five to thirty-five days prior to execution month, the scheduler will be finished compiling the set of approximately 250 aircraft missions. The resulting schedule is called the *initial cut*.

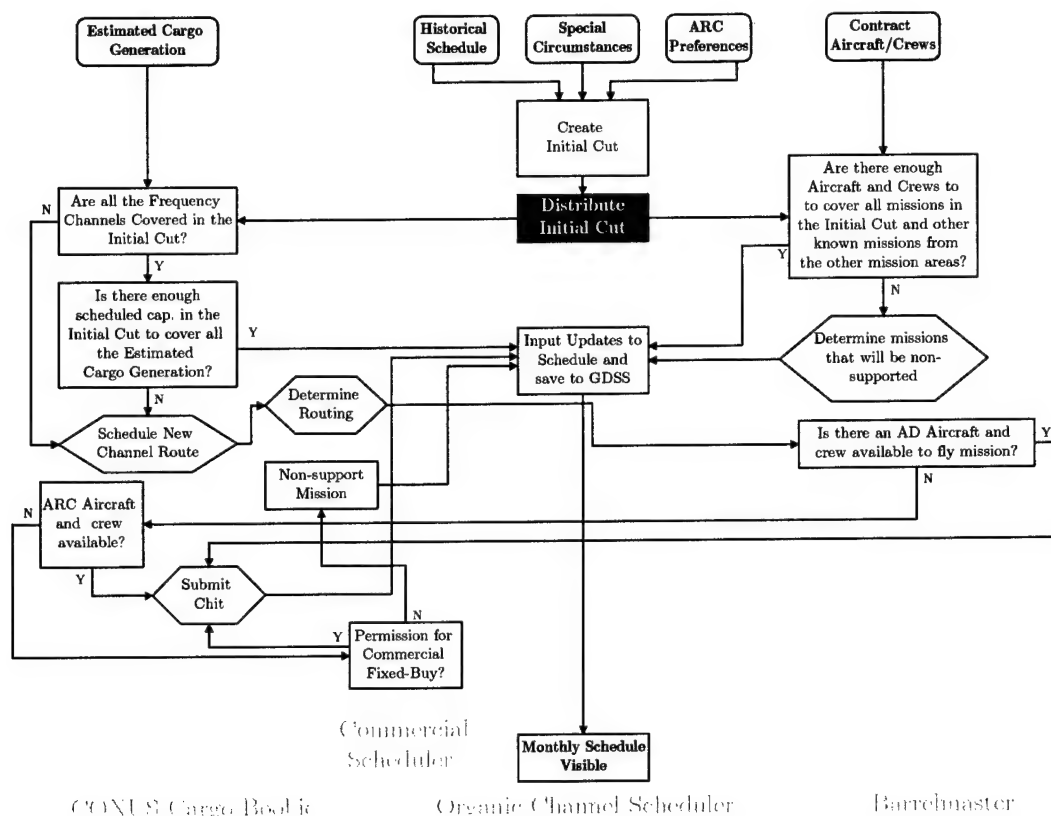


Figure 2-11: The current channel route scheduling process

Next, the organic channel scheduler passes the initial cut to the CONUS cargo bookie and the barrelnmaster for review. The CONUS cargo bookie requests changes to the schedule if there are not enough channel route missions to move all the estimated cargo demands, and the barrelnmaster requests changes based on aircraft availability. The organic channel scheduler incorporates these changes and saves the updated missions in the GDSS database. The schedule will be published in the *Channel Route Bulletin* (Bulletin) for all customers to view thirty days prior to execution month. The Bulletin is similar to the master flight listing many airlines publish. Specifically, it contains information on the operating times and locations of all channel route missions for the upcoming month. When changes to the set of channel route missions is required after the Bulletin has been published, the organic channel scheduler will update the channel route missions in GDSS. Changes range from a complete cancellation of a channel route mission to delaying takeoffs times or the day a channel route mission begins at the APOE. Although the missions are updated in GDSS, the Bulletin is not republished to reflect these changes.

Essentially, the organic channel scheduler performs two tasks. First, when creating the initial cut, the organic channel scheduler must manually pick the best set of channel route missions for the execution month. The definition of "best" is arbitrary and, as a result, the organic channel scheduler must resort to using the historical initial cuts and feedback from the CONUS cargo bookie and barrelmaster. Second, the organic channel scheduler makes decisions about channel route mission modifications throughout the planning process and the execution month. For example, the scheduler may need to change the takeoff time at one stop in the channel route mission to adhere to WMOG and PMOG limitations, or change the flying wing that is responsible for executing the mission because of aircraft availability.

When creating the initial cut, the organic channel route scheduler does not consider the amount of cargo expected to be in the channel route network. As long as the amount of the cargo does not change from month to month, this technique is adequate. However, the level of cargo is not necessarily constant from month-to-month. This is important because the expected cargo flow in the channel route network plays a major role in determining how aircraft flows through the channel route network. Similarly, the number of available aircraft will have a significant impact on the flow of aircraft. The organic channel route scheduler incorporates cargo and aircraft availability information *after* the initial cut has been distributed to the cargo bookie and the barrelmaster. Incorporating this information while creating the initial cut would eliminate the need for an entire iteration in the scheduling process. Furthermore, the channel route scheduler must rely on intuition and experience to evaluate the potential system-wide influences of changing a single channel route mission. This is monumental task considering there are approximately 250 missions that span an entire month. Additionally, the organic channel route scheduler does not explicitly consider the FHP during the planning process. Although there are monthly FHP target levels for each aircraft, these are only considered when directives from higher levels are passed down instructing the organic channel route scheduler to schedule a specific aircraft type more or less frequently.

Unfortunately, the organic channel route scheduler does not keep track of the changes made to the initial cut. Keeping track of these changes would allow the organic channel route scheduler to compare the final schedule (with all the changes) to the initial cut and determine any trends in the changes made to the initial cut from month-to-month. Although the changes are specific to each month, considering trends might be valuable when creating the initial cut.

2.3.3.3 Commercial Scheduler

The commercial scheduler is responsible for expansion-buy contracts (fixed-buy contracts are signed for an entire fiscal year, which means they operate without oversight from the

commercial scheduler). The need for commercial augmentation comes about when the barrelmaster non-supports a mission, or when the CONUS cargo bookie adds a new channel route mission and there is no organic airlift available. The CONUS cargo bookie must submit an expansion-buy request to the commercial scheduler. This requires the approval from the Global Channels Directorate Commander (a Colonel), the TACC commander (a Major General), and, when the ARA is at a certain level, the AMC DO (a Major General). This approval process is required because expansion-buy contracts are short notice for the commercial carriers and, as result, extremely expensive for AMC.

The approved expansion-buy request is then input to the *Commercial Operations Integrated System* (COINS) by the commercial scheduler. The commercial scheduler will specify the route to be flown, the starting day, the tonnage to be moved, and the mileage of the route. The value of the contract is determined by the tonnage, the mileage, and the agreed upon expansion-buy rate. After the contracts division at AMC has reviewed the request within COINS for legalities, the request is posted to the COINS *bulletin board* for all commercial carriers to view. If a commercial carrier is interested in the expansion-buy contract and has the resources to serve the contract, they will notify the commercial scheduler and the contract will be awarded based on the distribution of peacetime bid points among the interested commercial carriers.

After the contract is awarded the commercial scheduler will work with the cargo bookie, the organic scheduler, and the commercial carrier to set specific timing for the mission such that the mission does not violate WMOG restrictions and operating hours at each aerial port along the route.

2.3.3.4 Cargo Bookies

There are two groups of cargo bookies involved in the channel route planning process, the CONUS cargo bookies and the offshore cargo bookies. The CONUS cargo bookies are responsible for all cargo movements that originate in the CONUS and are destined for OCONUS locations. The offshore cargo bookies are responsible for cargo that originates at OCONUS locations and is destined either for another OCONUS location or a CONUS location.

CONUS CARGO BOOKIE

Similar to the organic channel schedulers, there are a total of three CONUS cargo bookies. At any given time, two are in the planning period and one is managing the execution months' schedule (refer to *Figure 2-9* in §2.3.3).

Cargo Generation Estimates

At the beginning of the planning period, the CONUS cargo bookie estimates the amount of cargo expected to be at each of the CONUS APOE's. This yields an estimate for the amount of cargo expected for each validated requirements channel on each day of the month. This daily estimate is referred to as the *cargo generation estimate*. During this process the cargo priority (e.g., triple-9, TP-1) is not considered. The CONUS cargo bookie assumes that the APOE's will consider cargo priority when loading the aircraft. Another assumption made by the CONUS cargo bookie is that the APOE's will correctly load the aircraft to coincide with the aircraft route. For example, cargo destined for the first stop should be placed at the rear of the aircraft so it can be easily offloaded without unloading the entire aircraft. Furthermore, the cargo generation estimate is only for bulk cargo. Other categories of cargo are handled on a case-by-case basis.

Currently, there is no standard method by which all CONUS cargo bookies calculate the cargo generation estimate. For example, some CONUS cargo bookies use an average of actual cargo generation in the previous three months, while others use the previous three months and the same month from the previous year. General rules of thumb and experience play a large role in how a CONUS cargo bookie calculates the cargo generation estimate.

Furthermore, CONUS cargo bookies do not analyze the accuracy of their cargo generation estimates after the actual cargo levels have been realized. Simple analysis such as plotting the cargo generation estimates against the actual cargo levels after the execution month would provide insight into the ability of the current estimation techniques to accurately forecast cargo generation. It is likely that forecasting techniques (e.g., exponential smoothing and ARIMA models) could improve the cargo generation estimates. Specifically, these techniques could identify the presence of weekly, monthly, and seasonal trends. Greater visibility of the cargo would also help the CONUS cargo bookie calculate cargo generation estimates. Currently, customers do not submit accurate forecasts of their expected demand. Similarly, during the execution month, the CONUS cargo bookie does not have visibility of the cargo being shipped to an APOE until it actually arrives. This forces the CONUS cargo bookie to *react* to the cargo in the channel route network. Reacting, rather than anticipating, leads to sub-optimal use of resources, poor customer service, and does not allow cargo to flow seamlessly through the military logistics system.

The organic channel scheduler gives the initial cut to the CONUS cargo bookie between 35 and 45 days prior to the start of the execution month. The CONUS cargo bookie compares the *Sequence Listing* and the cargo generation estimate to the initial cut to determine if there are enough scheduled channel route missions to cover all the frequency channels and all the

estimated cargo generation. If all the frequency channels are covered and all the estimated cargo generation is covered, the CONUS cargo bookie makes no changes to the initial cut.

Add-on and Modified Channel Route Missions during the Planning Period

If there are not enough channel route missions in the initial cut to cover all the frequency channels and/or all the estimated cargo generation the CONUS cargo bookie will submit a *chit* to the organic channel scheduler. A *chit* is a request to add a new channel route mission or modify an existing channel route mission. To submit a *chit*, the CONUS cargo bookie will determine the aircraft route, starting day, and aircraft type to fly the channel route mission. The cargo bookie must confirm with the barrelmaster that an organic aircraft is available to fly the channel route mission. If an organic aircraft is not available the CONUS cargo bookie will contact the ARC component squadrons and determine if they will fly the mission. If not, the CONUS cargo bookie will seek permission for commercial augmentation via an expansion-buy contract.

After the CONUS cargo bookie has confirmed that an aircraft is available, the *chit* is submitted to the organic channel scheduler. The organic channel scheduler will determine the specific timing (i.e., take-off times, crew rest times) and operating characteristics (i.e., onload, offload, crew rest), input the new channel route mission via CAMPS, and save the mission in the GDSS database.

Add-on and Modified Channel Route Missions during the Execution Month

The process of adding and/or amending a channel route mission during the execution month is slightly different. During the execution month there are several situations that warrant changes to be made by the CONUS cargo bookie. For instance, a higher priority mission from another mission area can cause a channel route mission to be non-supported or significantly modified. This would leave the CONUS cargo bookie with backlogged cargo at one or more APOE's. Another example is if the actual cargo generation is higher than the cargo generation estimate. The CONUS cargo bookie monitors the actual daily levels of cargo and the *station workload summary* at each CONUS APOE. The station workload summary lists all the channel route missions scheduled to visit the aerial port within a 72-hour window. The CONUS cargo bookie uses the station workload summary to ensure that there are enough scheduled channel route missions to move the actual daily levels of cargo at each CONUS aerial port. If there are not enough channel route missions to move all the cargo, the CONUS cargo bookie must submit a *chit* to the organic channel scheduler.

The process of obtaining commercial augmentation varies slightly from the process followed during the planning period. To obtain permission for commercial augmentation, the CONUS cargo bookie must confirm that there are neither active organic aircraft nor ARC aircraft available to fly the mission and the cargo must have a port hold time (i.e., been sitting at the APOE) between 48 and 72 hours. The "48 to 72 hour rule" ensures that the top-dollar expansion-buy contracts are used only as a last resort.

During the execution month, there are also opportunities for a CONUS cargo bookie to schedule channel cargo to be moved using *opportune missions*, which either opens capacity on aircraft flying channel route missions or alleviates the need for a scheduled channel route mission. An opportune mission occurs when an aircraft operating in another mission area (i.e., SAAM, Exercise) has residual capacity and is traveling between the origin and destination of a channel. Opportune missions do not influence the development of the initial cut because they become visible to the organic scheduler and CONUS cargo bookie only hours prior to the time they are available to transport cargo.

End-of-Month Report

At the end of each execution month the CONUS cargo bookie submits an end-of-month performance report. Included in this report are summary statistics capturing the number of channel route missions flown, non-supported, added, and modified. Additionally, the total *CONUS outbound utilization rate* and the daily utilization rate for each CONUS APOE are reported. The daily CONUS outbound utilization rate, u , for day d at CONUS APOE b is defined in equation (2.1). It is the ratio of the sum of all the actual cargo generation on day d at APOE b to the sum of the capacities of all the aircraft scheduled to pass through the APOE b on day d .

$$u_d^b = \frac{\sum_{r \in R} g_{d,r}^b}{\sum_{a \in A} c_a}, \text{ where} \quad (2.1)$$

- A is the set of all aircraft scheduled to visit CONUS APOE b on day d ,
- R is the set of all requirements channels such that the origin of r is CONUS APOE b ,
- c_a is the TWCF planning weight of aircraft a ,
- $g_{d,r}^b$ is the actual cargo generation for channel r at APOE b on day d .

Let D be the set of all days in the execution month. The total utilization rate, U^b , at APOE b is:

$$U^b = \frac{\sum_{d \in D} u_d^b}{|D|}. \quad (2.2)$$

It should be noted that this metric only measures the first leg of each route; namely, the CONUS to OCONUS flight leg. In general, there is much more cargo shipped from CONUS to OCONUS than there is within OCONUS. As a result, aircraft operating channel route missions that originate at CONUS APOE's are relatively full on the first flight leg. Consequently, the CONUS outbound utilization rates defined in equations (2.1) and (2.2) are inflated relative to the aircraft utilization over the entire channel route mission. The utilization rates defined in equations (2.1) and (2.2) can be greater than one because the TWCF planning weights are used, opposed to the maximum aircraft capacity.

Ultimately, this metric drives the planning behavior and tendencies of the CONUS cargo bookies. As a result, on a daily basis, not much consideration is given to the utilization rate over all the flight legs in a channel route mission. Different performance metrics might demonstrate that a lower utilization rate on the first leg of a channel route mission corresponds to a higher utilization rate over the entire channel route mission and/or higher utilization rates on other channel route missions.

OFFSHORE CARGO BOOKIE

The offshore cargo bookie's job is drastically different than the CONUS cargo bookie. The offshore cargo bookie is not involved in creating the initial cut. In fact, the offshore cargo bookie only monitors the amount of cargo at the OCONUS aerial ports each day to ensure that a sufficient number of channel route missions are scheduled through each aerial port to transport the cargo to its destination. The offshore cargo bookie has limited options to deal with aircraft capacity shortages. Working with the organic scheduler, the offshore cargo bookie can attempt to add a flight leg to an existing channel route mission. The offshore cargo bookie does not have the option to submit a chit to create a new channel route mission solely to move OCONUS cargo. The only way the offshore cargo bookie can add a mission is if the CONUS cargo bookie also has a shortage of aircraft capacity. In this case the offshore bookie will coordinate with the CONUS cargo bookie to ensure the aircraft routing submitted by the CONUS cargo bookie includes the OCONUS aerial ports that have a shortage of aircraft

capacity. Unlike the CONUS cargo bookie, the offshore cargo bookie does not submit an end-of-month report.

2.3.3.5 The Barrelmaster

The barrelmaster is responsible for allocating aircraft and aircrews among four of the five mission areas: contingencies, exercises, SAAMs, and channel routes. Aircraft allocation for JA/ATT's is handled by the individual aircraft squadrons. Additionally, the barrelmaster is responsible for determining the locations to preposition aircrews.

Approximately 30 to 45 days prior to the start of the execution month each of the aircraft wings will submit the number of contract aircraft and aircrews to the barrelmaster. This is an upper limit on the number of aircraft and aircrews that can be assigned to the four mission areas. The number of contract aircraft and aircrews changes from month-to-month, and can fluctuate daily.

The main tool used by the barrelmaster is the *commitment matrix*. The commitment matrix is a spreadsheet that contains information about the contract number of aircraft and aircrews by aircraft wing and the number of missions (from all mission areas) scheduled, for each aircraft type. The spreadsheet is integrated with GDSS and is automatically updated to reflect all missions that are in the GDSS database. As long as the number of scheduled missions is less than the contract numbers, the barrelmaster is not required to make changes. However, if the scheduled missions use more than the contract number of aircraft or aircrews, the barrelmaster must work with the planners from each of the mission areas to resolve the conflict.

There are several ways that a barrelmaster can resolve conflicts. One method is to drop one or more legs of a channel route mission. Another method is to change the aircraft type operating the channel route mission. Both methods change the aircraft capacity over the channel route mission. This might require the CONUS cargo bookie to find other means to move cargo located at the affected aerial ports. If these remedies do not resolve the conflict, the barrelmaster must ultimately non-support the lowest priority mission(s), which leaves the CONUS cargo bookie to find airlift capacity.

If the barrelmaster must non-support a channel route mission it is often the case that several channel route missions are candidates to be non-supported. Currently, the barrelmaster must use experience and feedback from the CONUS cargo bookie and the organic channel route scheduler to minimize the impact on the entire channel route network. Other than experience, there is nothing to help the channel route planners analyze the network-wide effect of either dropping a channel route mission, offering it to the ARC component, or offering it to commercial carriers.

The barrelmaster is also responsible for determining when and where to preposition aircrews. Currently, the barrelmaster must rely on intuition and experience to guide this determination. Although the set of aerial ports that are candidates for prepositioned aircrews is limited, the barrelmaster is only able to evaluate a subset of the potential options to preposition aircrews. Additionally, it is difficult for the barrelmaster to determine the influence that different prepositioning strategies will have on the entire channel route network.

When the barrelmaster is forced to modify the existing set of scheduled missions in GDSS, issues such as aircraft reliability, the servicing facilities of the stops along the route, and the geographic location of the aircraft must be considered. Depending on the priority of the mission, the barrelmaster will not assign an aircraft type that historically has low reliability. Furthermore, there are certain locations that are not well equipped to handle maintenance issues of certain aircraft types. The barrelmaster attempts to keep those aircraft types from those locations, especially on high priority missions. Another issue that is considered is the aircraft's home location relative to the origin APOE. The customer must pay for positioning and depositioning legs. If able, the barrelmaster attempts to minimize these costs.

The barrelmaster must also deal directly with the various aircraft squadrons. There are instances when aircraft squadrons would prefer to delay the start of a mission anywhere from a few hours to several days. These aircraft squadrons often call the organic channel schedulers to resolve these issues. The organic channel scheduler should, in turn, contact the barrelmaster to ensure that timing changes will not cause conflicts with aircraft availability for other mission areas. This step in the process could be avoided if the aircraft squadrons contacted the barrelmaster directly.

2.3.3.6 The Floor

Within 24 hours of execution, all responsibility for a channel route mission is transferred to the *floor*. The floor consists of the *Aerial Port Control Center (APCC)* and the *cells*. The APCC performs the role of the cargo bookies, assuming responsibility for all cargo on the aircraft and, to a limited extent, managing the cargo at the aerial ports along the channel route mission. The cells perform the role of the organic channel route scheduler and barrelmaster, assuming responsibility for the aircraft and routing. Both the APCC and the cells communicate with the barrelmaster to coordinate the effects of mission changes.

The APCC maintains a detailed report of all the events that take place during the execution of a channel route mission. This minute-by-minute report is called the *flight following report*. The flight following report details any delays and how the delays will influence the remainder of the flight, with an emphasis on how the flight delays will influence the cargo

movements planned for the flight. For example, if an aircraft is delayed for maintenance, the APCC will determine how to fit the cargo from the delayed aircraft on an aircraft scheduled to come through the aerial port. The APCC will also monitor the cargo levels at all the aerial ports to ensure that high priority cargo is not backlogged. In some situations, the APCC will divert an aircraft to an unscheduled aerial port to move urgent cargo.

2.4 Motivation for Models

The purpose of this section is to revisit AMC's three broad operational objectives in the context of the channel route planning process. The focus is to highlight how each of these objectives is taken into consideration during the channel route planning process and to motivate the models presented in the remaining chapters.

2.4.1 Readiness

The FHP is not explicitly considered by either the organic channel route scheduler, the CONUS cargo bookie, or the barrelmaster. The FHP is only taken into consideration when guidance is passed down from a higher level and when the fixed-buy analysis is performed. This allows AMC to make mid-course corrections. Considering the monthly FHP target levels while creating the initial cut might prevent the need for drastic mid-course corrections. Furthermore, an analytical approach would provide insights into how changes to the channel route schedule during the execution month influences AMC's ability to reach the target number of flying hours.

2.4.2 Customer Service

Customer service is the cargo bookie's main consideration during the planning period and especially during the execution month. During the planning period the CONUS cargo bookie ensures that all frequency channels are covered. Throughout the execution month, the CONUS cargo bookie and the offshore cargo bookie monitor the amount of cargo at each aerial port to ensure that all cargo is being moved. However, improving the cargo generation estimates and considering the cargo generation estimates during the creation of the initial cut has the potential to find new, non-obvious channel route missions and flow patterns for the cargo through the channel route network. Ultimately, this would mean that the channel route schedule would require fewer changes during the execution month and be less susceptible to the influence of higher priority mission areas. Fewer required changes to the channel route schedule has three potential benefits. First, AMC's customers would experience more predictable service.

Second, fewer expansion-buy missions would be required. Third, a more stable channel route network has a positive impact on readiness and NOR because a stable channel route network leads to a predictable source of flying hours and revenue. Additionally, considering the estimated cargo generation levels and aircraft availability during the creation of the initial cut would significantly reduce the amount of time required to perform an entire iteration of the planning process.

2.4.3 Net Operating Result

NOR is considered during the fixed-buy analysis and by the CONUS cargo bookie through the TWCF planning weights. Considering the estimated cargo generation levels during the creation of the initial cut might create a channel route schedule that allows cargo to flow through the channel route network in such a way that aircraft are utilized more efficiently. Currently, the CONUS cargo bookie pays close attention to this, but *only* on the first CONUS to OCONUS flight leg. Considering the TWCF planning rates on a network wide level might increase the worldwide utilization rate, which in turn would help keep NOR near zero.

2.5 Summary

This chapter has introduced channel routes in the context of all AMC operations. Furthermore, this chapter has familiarized the reader with the current channel route planning and execution process and the limitations of the current approach. These shortcomings motivate the necessity of an analytic approach to create a thirty-day channel route schedule. The formulations and models developed in the remaining chapters are initially aimed at improving the quality of the initial cut. Additionally, we will explore how the models and formulations used to improve the quality of the initial cut can be used to help the channel route planners in other capacities of the channel route planning process.

[This page intentionally left blank]

3 Functional Analysis & Traditional Modeling Approaches

The previous chapter presented the current channel route planning process from an operational perspective. The purpose of this chapter is to develop a network design formulation intended to support the development of the initial cut (see §2.3.3.1). This chapter begins with a functional analysis of the current channel route planning process. The purpose of the functional analysis is to identify the inputs and outputs relevant to the models and formulations presented in the remainder of this thesis.

Prior to developing the formulations, the physical channel route network is represented using a framework known as a *time-space graph*. Using the time-space graph, an *arc-based network design formulation* and a *path-based network design formulation* are developed. The two formulations are compared, which motivates the development of a decomposition strategy based on Dantzig-Wolfe decomposition [2],[15]. In the decomposition, a *master-level scheduler* determines the best set of cargo paths and aircraft routes, which are generated by two lower-level planners (i.e., subproblems). The two subproblems are shortest path problems that use dual prices to identify cargo path columns and aircraft route columns with negative reduced cost.

The decomposition is implemented using a sample data set. The results are presented and used to motivate the need for an alternate formulation with more desirable computational behavior.

3.1 Functional View of the Channel Route Scheduling Problem

The purpose of this section is to identify the inputs and outputs of the channel route planning process, without considering the specific operations that are used to create the desired outputs from the set of inputs (described in §2.3.3). In fact, the specific operations are viewed as a “black-box”, as illustrated in *Figure 3-1*. Currently, the “black-box” used at the TACC is a labor-intensive process, depicted in *Figure 2-10* (and repeated in *Figure 3-3a*). The purpose of this research is to explore mathematical models that will improve the initial cut, using the same set of inputs currently used in the channel route planning process. An improved initial cut requires fewer aircraft and/or takes less time for the channel route planners to create. Ideally, an improved initial cut will also better capture AMC’s three operational objectives – *readiness*, *customer service*, and *NOR* – introduced in §2.1.3.

3.2 An Improved Operational View

In the previous chapter, we highlighted several shortcomings of the current channel route planning process. The purpose of this section is to illustrate how – from an operational perspective – these challenges can be addressed using a *mixed integer program* (MIP).

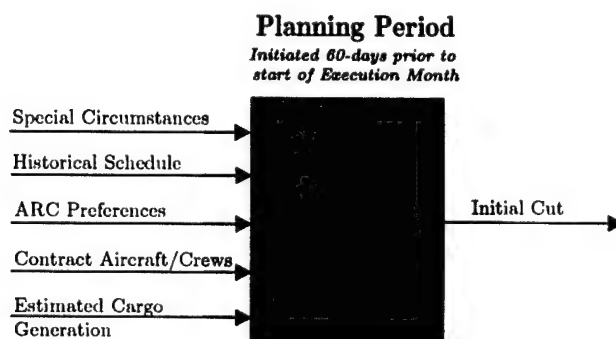


Figure 3-1: Functional View of the Channel Route Planning Process

Developing a mathematical model of a large-scale, complex system such as the channel route network is difficult because there are many highly related decisions and integrated processes. Ideally, the entire system should be modeled to account for all the interactions in the system. However, even if formulating such a model is possible, it is often not tractable. Dealing with tractability issues requires that different formulation strategies and/or different solution techniques be utilized. A common reformulation strategy is to divide a large, intractable problem into a series of smaller, tractable subproblems, which are solved sequentially. Although disaggregating the problem will typically yield a sub-optimal solution for the entire system, the goal is to get a reasonable solution that is feasible. It is not uncommon that the subproblems must be disaggregated further in order to obtain tractable formulations. Each time a problem is divided into more subproblems, a tradeoff is made between tractability and optimality. Thus, the challenge is to find a way to split the problem into a set of smaller, more tractable problems while capturing as many interactions as possible.

An example from the literature is the airline schedule planning problem. *Figure 3-2* illustrates the decomposition of the airline schedule planning problem [12]. The *schedule generation problem* is concerned with generating a feasible schedule of flights (i.e., the cities to service and the times to service them). The *fleet assignment problem* assigns fleet types (i.e., Boeing 737) to each flight in the schedule created during schedule generation. Given the fleet assignments, the *maintenance routing problem* determines the specific aircraft (tail numbers) that will operate each flight. Finally, the *crew scheduling problem* assigns aircraft crews to the flights in the schedule created during maintenance routing. Each of these problems is addressed individually in the literature [6],[9],[34]. Additionally, capitalizing on recent advances in modeling and computing power, the literature addresses the benefits of integrating several of these subproblems. For example, Barnhart et al. [10] discuss the benefits of integrating the fleet assignment problem and the crew scheduling problem, and they propose a fleet assignment model that includes an approximation of the crew pairing problem. Similarly, Cohn and Barnhart [23] develop a crew pairing model that integrates key maintenance routing decisions.




Figure 3-2: The airline schedule planning problem

Although the airline schedule planning process is decomposed into several smaller subproblems, all the relevant data for each subproblem is considered *simultaneously* in order to optimize the subproblem. For example, the fleet assignment problem assigns the fleet types to flights using the estimated passenger demands because the number of passengers has a direct impact on the size of aircraft that is assigned.

In the current channel route planning process, the major drawback is that cargo generation estimates and aircraft availability information are not used to create the initial cut. Incorporating this information during the creation of the initial cut has the potential to yield significant improvements. Incorporating cargo generation estimates and aircraft availability information would change the current planning process (depicted in *Figure 3-3a*) to resemble the process depicted in *Figure 3-3b*.

It is important to note that the functional diagram in *Figure 3-1* still applies to *Figure 3-3b*. In other words, the same set of inputs currently used in the channel route planning process is utilized in *Figure 3-3b*, and the output from *Figure 3-3b* is the same as the output created from the current channel route planning process. As *Figure 3-3b* illustrates, the major improvement over the current channel route planning process (see *Figure 3-3a*) is that the optimization model used to create the initial cut is able to consider simultaneously, rather than sequentially, all of the relevant inputs when creating an initial cut.

3.3 Static Graph Representation

Applying analytic techniques to evaluate the channel route network requires that the physical system be mapped to the proper theoretical framework. Large-scale transportation problems, whose physical structure is similar to the channel route network, are often represented using a directed graph, $G(N,A)$, consisting of a set of nodes, N , and a set of arcs, A . There are many ways to translate the physical characteristics of the channel route network into such a graph.

One possible directed graph representation of the channel route network consists of nodes that correspond to each physical location at an aerial port, and arcs between each pair of nodes representing the movement (i.e., flow) of either cargo or aircraft between the two physical locations (see *Figure 3-4a*). The arcs between nodes indicate the direction of the flow between the two nodes; the tail of the node corresponds to the origin and the head of the node corresponds to the destination.

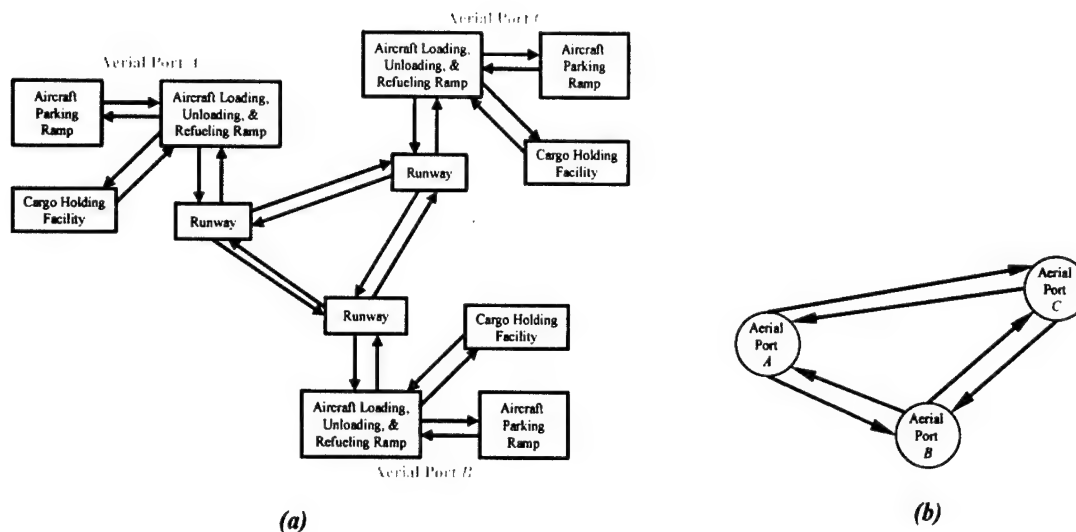


Figure 3-4: Graph Representations of the Channel Route Network

A different directed graph representation of the channel route network consists of nodes that correspond to an entire aerial port, and directed arcs between two nodes representing the flow of either cargo or aircraft, or both, between the two aerial ports (see *Figure 3-4b*).

The level of fidelity captured in a graph has a significant impact on the size, and hence, the tractability of the mathematical formulation that is used to describe the graph. As the scope of this research is to model the flow of cargo and aircraft between aerial ports, it is not necessary to model all the cargo and aircraft activities that take place at each aerial port. Consequently, the graph in *Figure 3-4b* provides the level of detail required to successfully capture the channel route network of interest.

3.4 Time-Space Graph Representation

This section introduces the general concept of a time-space graph, and an extension of this concept that reduces the number of variables and constraints that must be explicitly included in a mathematical formulation.

3.4.1 General Time-Space Graph

The graph in *Figure 3-4b* is static, which means it allows neither cargo flows nor aircraft flows to be modeled in the time dimension. The static graph in *Figure 3-4b* is modified to create a *time-space graph*, which captures cargo flows and aircraft flows through both the

spatial and temporal dimensions. The time-space graph is constructed by creating a node for each aerial port during every time period in the *planning horizon*. The planning horizon is the entire time period of interest. For example, the planning horizon for the channel route planning problem is thirty days. Each node in the time-space graph is indexed by an aerial port and a time period. For instance, the index (A,t) corresponds to aerial port A in time t . Directed arcs connect the nodes and represent cargo and aircraft flows in time and space. A time-space graph representation of the static graph in *Figure 3-4b* is illustrated in *Figure 3-5*.

An arc in the time-space graph that connects two nodes with the same aerial port index in consecutive time periods is a *ground arc* (represented by a dashed arc in *Figure 3-5*). Flow over a ground arc indicates that cargo and/or aircraft remains at the aerial port from time t to time $t+1$. All other arcs are *flight arcs* (represented by solid arcs in *Figure 3-5*).

The formulations in the remainder of this thesis use a time-space graph to create a *channel route schedule* (schedule). A *schedule* contains routing information for the cargo flows, and routing and execution details for the aircraft flows. The execution details describe the aircraft action at a particular aerial port. Specifically, an aircraft can either *onload* or *offload* cargo (or both) at an aerial port. Additionally, an aerial port can serve as either the start of a *positioning* flight leg or the end of a *depositioning* flight leg. A positioning flight leg occurs when an empty aircraft flies from its home base to the origin of a channel route. Similarly, a depositioning flight leg occurs when an empty aircraft returns to its home base from the destination of a channel route. It is important to note that the schedule does not include precise timing, but rather an initial sense of timing. *Example 3-1* illustrates the flow of cargo items and aircraft through the time-space graph and a resulting schedule.




Figure 3-5: Time-Space Graph of the Channel Route Network

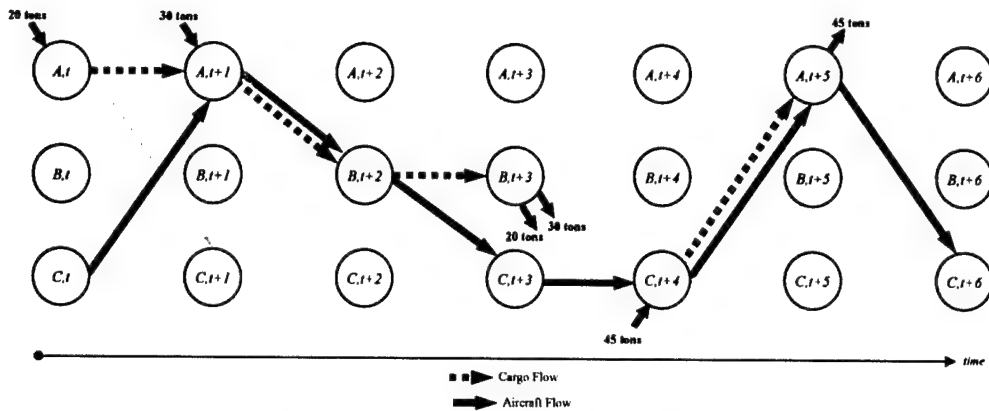


Figure 3-6: Aircraft and cargo flow through the time-space graph

Example 3-1: Consider the time-space graph (three aerial ports – A,B,C – during seven time periods) depicted in Figure 3-6, and one aircraft with 50 tons of capacity that is located at node (C,t) . Furthermore, the following three requirements channels must be served:

Channel 1: 20 tons from node (A,t) to node $(B,t+3)$

Channel 2: 30 tons from node $(A,t+1)$ to node $(B,t+3)$

Channel 3: 45 tons from node $(C,t+4)$ to node $(A,t+5)$.

To service the three channels, the aircraft begins by traversing the flight arc from node (C,t) to node $(A,t+1)$. Simultaneously, the channel 1 cargo remains at aerial port A, flowing on the ground arc from node (A,t) to node $(A,t+1)$. The channel 1 and channel 2 cargo is onloaded to the aircraft at node $(A,t+1)$ and the aircraft departs on the flight arc from node $(A,t+1)$ to node $(B,t+2)$. At node $(B,t+2)$, all the cargo is offloaded and flows over the ground arc between node $(B,t+2)$ and its destination, node $(B,t+3)$. The aircraft continues on the flight arc between node $(B,t+2)$ and node $(C,t+3)$. The aircraft remains at aerial port C for one time period, flowing on the ground arc from node $(C,t+3)$ to node $(C,t+4)$. At node $(C,t+4)$, the channel 3 cargo is onloaded to the aircraft and delivered to its destination via the flight arc from node $(C,t+4)$ to node $(A,t+5)$. The channel 3 cargo is offloaded, and the aircraft returns to aerial port C via the flight arc from node $(A,t+5)$ to node $(C,t+6)$. Let the notation $(\text{node } i) \rightarrow (\text{node } j)$ represent movement over the arc originating at node i and terminating at node j . The schedules for the channel cargo and the aircraft are represented as:

CHANNEL 1: $(A,t) \rightarrow (A,t+1) \rightarrow (B,t+2) \rightarrow (B,t+3)$

CHANNEL 2: $(A,t+1) \rightarrow (B,t+2) \rightarrow (B,t+3)$

CHANNEL 3: $(C,t+4) \rightarrow (A,t+5)$

AIRCRAFT: $(C,t) \rightarrow (A,t+1) \rightarrow (B,t+2) \rightarrow (C,t+3) \rightarrow (C,t+4) \rightarrow (A,t+5) \rightarrow (C,t+6)$, or

<i>Aerial Port</i>	<i>Time</i>	<i>Execution Details</i>
<i>C</i>	<i>t</i>	<i>start of positioning flight leg</i>
<i>A</i>	<i>t+1</i>	<i>onload Channel 1 and Channel 2 cargo</i>
<i>B</i>	<i>t+2</i>	<i>offload Channel 1 and Channel 2 cargo</i>
<i>C</i>	<i>t+3</i>	<i>end of depositioning flight leg</i>
<i>C</i>	<i>t+4</i>	<i>onload Channel 3 cargo</i>
<i>A</i>	<i>t+5</i>	<i>offload Channel 3 cargo</i>
<i>C</i>	<i>t+6</i>	<i>end of depositioning flight leg</i>

3.4.2 Incorporating Constraints in the Time-Space Graph

The previous section introduced the general concept of a time-space graph, which included an arc between *each* pair of nodes. However, it is neither necessary nor desirable to include every arc in the time-space graph. Excluding arcs that cannot be used in the solution is important because fewer constraints and variables are required to mathematically represent the time-space graph. It is known a priori that channel route missions are subject to the routing rules presented in §2.1.3.1, that each aircraft type has a maximum flying range, and that certain service levels are desired. Building these constraints into the graph structure prevents the need to explicitly define the constraints in a mathematical formulation. The remainder of this section provides an example of how constraints are incorporated for the notional channel route network described in *Example 3-2*.

Example 3-2: Consider five aerial ports, *A*, *B*, *C*, *D*, and *E*. Only aerial port *A* is a squadron aerial port. The notation (i,j,t) indicates that a validated channel exists between aerial port *i* and aerial port *j* and that the cargo's available to load time is *t*. The available to load time (the time period the cargo is available at its origin aerial port *i* for pickup) is denoted as ALT^i . The time *t* is specified as a single time period for requirements channels and as a frequency of service for frequency channels. The following six channels must be served: $(A,B,0)$, $(A,C,0)$, $(A,C,1)$, $(A,E,1/\text{week})$, $(B,A,1)$, and $(C,D,2)$. For simplicity, it is assumed that the flight time between each aerial port is one time period, and that all the available aircraft have the range capabilities to fly between any two aerial ports. T_{ij} is the flight time between aerial port *i* and aerial port *j*. Additionally, let the maximum port hold time (PHT) be 2 time periods. Finally, if an aircraft departs its squadron aerial port (aerial port *A*, in this example) it must return to its squadron aerial port prior to the end of the planning horizon.

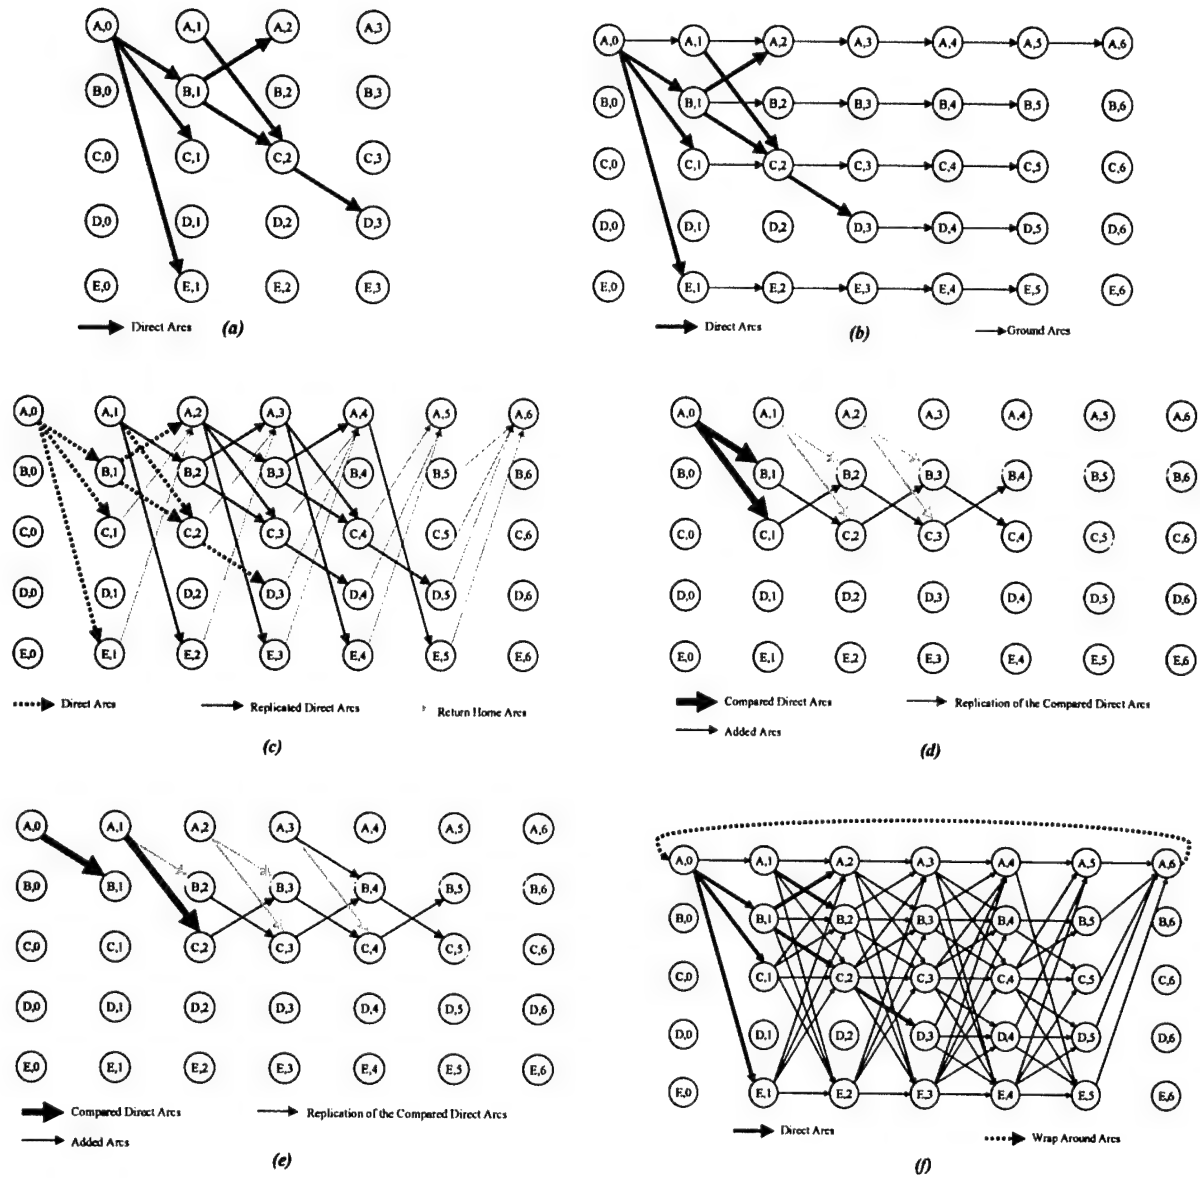


Figure 3-7: Construction of a time-space graph for a notional channel route network

To construct the improved time-space graph, we search the list of requirements channels for the maximum available to load time (ALT^{max}), which is $t=2$ in *Example 3-2* (from channel (C,D,2)). The time required to travel directly from the origin to the destination of the requirement containing ALT^{max} is noted as T^{max} . A node is created for each aerial port, and is replicated for each time period from $t=0$ through $t=ALT^{max}+T^{max}$. An arc is created between the origin node and the destination node for each requirements channel. The arc begins at the origin node i in time $t = ALT^i$ and terminates at the destination node j in time period

$t = ALT^i + T_{ij}$. For each frequency channel, an arc is added from the origin i to the destination j of the frequency channel, starting in the first time period of the frequency of service interval (e.g., for a once per week frequency of service, an arc is added that begins in the first time period of that week). The length of the arc is T_{ij} . These initial arcs are referred to as the *direct arcs*, and represent direct travel from the origin to the destination of each channel. These steps are illustrated in *Figure 3-7a*.

Next, additional nodes are added to the time-space graph to allow the channel with ALT^{max} to remain at its origin for PHT time periods. A final time period is added so that each aircraft is able to return to its squadron aerial port prior to the end of the planning horizon. At squadron aerial port nodes, we add ground arcs for all time periods after the earliest ALT for that aerial port. At non-squadron aerial port nodes, we add ground arcs starting in the time period of the first ALT at the aerial port and ending one time period prior to the final time period in the planning horizon. Omitting a ground arc between the last two nodes of a non-squadron aerial port prevents aircraft from remaining at a non-squadron aerial port. These steps are illustrated in *Figure 3-7b*.

Next, each of the direct arcs corresponding to a requirements channel is replicated for PHT time periods, capturing the fact that cargo can remain at its origin for at most PHT time periods. These arcs build the service level constraints into the structure of the graph because every feasible solution is guaranteed to satisfy the maximum PHT restriction (e.g., two days). The direct arcs corresponding to the frequency requirements are replicated for each time period in the frequency of service time interval (e.g., a once per week frequency of service would require that the direct arc be replicated for each time period in the one week time interval). These arcs are referred to as *replicated direct arcs*. Additionally, for each non-squadron aerial port node i that has either an incoming direct arc, a replicated direct arc, or a ground arc, an arc is added from that node back to each squadron aerial port node j . These arcs are referred to as *return home arcs* and have a length of T_{ij} . *Figure 3-7c* illustrates these modifications.

The graph in *Figure 3-7c* captures all the possible direct routes serving requirements channels and frequency channels. A time-space graph that only models direct routes doesn't, however, allow aircraft routes that service multiple channels. For instance, consider the two channels $(A,C,0)$ and $(A,B,0)$. Unless a direct flight arc exists from node $(C,1)$ to node $(B,2)$, the route $A \rightarrow C \rightarrow B \rightarrow A$ cannot be modeled. To capture these types of routes, each channel is compared to other channels of the same type (i.e., requirements channel or frequency channel) that have the same origin and that have ALT's within the maximum PHT of one another. An arc is added during each time period between the destinations of two channels being compared, if one does not exist. Two examples are shown in *Figure 3-7d* (two channels – $(A,B,0)$ and

$(A,C,0)$ – with the same ALT) and *Figure 3-7e* (two channels – $(A,B,0)$ and $(A,C,1)$ – with different ALT's).

Comparing two channels of different types is slightly different. Each frequency channel is compared to all other requirements channels whose origin is the same as the frequency channel's origin. For example, given the frequency channel $(A,E,1/week)$, the requirements channels $(A,B,0)$ and $(A,C,0)$ are flagged to be examined. Given one of the flagged requirements channels, i , if the destination of i is the origin of a requirements channel j that is not in the list of flagged requirements channels, then the frequency channel is compared to requirements channel j . In the case where $i=(A,C,0)$, requirements channel $j=(C,D,2)$ is chosen and compared with $(A,E,1/week)$. An arc between the two destinations, D and E, is added. The purpose of adding these arcs is to model the route $A \rightarrow C \rightarrow D \rightarrow E$, which services requirements channels $(A,C,0)$ and $(C,D,2)$, and frequency channel $(A,E,1/week)$. Adding these arcs will, however, allow a few missions that violate the routing rules presented in section 2.1.3.1. For instance, the route $A \rightarrow E \rightarrow D$ is feasible in the resulting time-space graph. This mission is not feasible, however, in the real-world system because there is neither a validated channel between A and D nor a validated channel between E and D. Although a few illegal routes are allowed, the total number of illegal routes, arcs, and nodes that are omitted is worth the effort of constructing the time-space graph as described.

Finally, wrap-around arcs – which are referred to as overnight ground arcs in the fleet assignment literature ([12],[34]) – are added at the squadron aerial ports to enable the aircraft balance of flow constraints. The wrap-around arcs originate during the last time period and terminate in the first time period. The final time-space graph is illustrated in *Figure 3-7f*. Note that many arcs have been excluded and various nodes (e.g., $(B,0)$) can be removed because they will never be used. For the remainder of the thesis, the term *graph* refers to a time-space graph, such as the one in *Figure 3-7f*.

3.5 Problem Classification

At its core, the channel route planning problem is a *network design problem*, which inherently includes a *multicommodity network flow problem*. The purpose of this section is to introduce two *multicommodity network flow formulations* and a *network design formulation* that are used to determine flow on general networks and that we will apply to the graph described above.

3.5.1 Multicommodity Network Flow Formulations

The following sections describe two *multicommodity network flow* (MCNF) formulations. MCNF formulations are used to determine the optimal flow of a set of non-homogeneous commodities through a graph, subject to set of constraints. The *arc flow formulation* is concerned with assigning flows to individual arcs in the graph. The *path flow formulation*, on the other hand, is concerned with assigning flows to *paths* through the graph.

3.5.1.1 Arc Flow Formulation

The objective of the arc flow formulation is to determine the least cost way to assign flow to individual arcs in the graph, subject to a set of constraints. The arc flow formulation is based on a directed graph $G(N, A)$, consisting of a set of arcs $(i, j) \in A$ and a set of nodes $i \in N$. Let $k \in K$ be the set of origin-destination commodity pairs. The linear cost per unit flow for commodity k on arc (i, j) is c_{ij}^k , and the capacity of arc (i, j) for commodity k is u_{ij}^k . Finally, b^k is the demand/supply for commodity k . The origin node i of commodity k , noted $O(k)$, is called a *supply* node. At $O(k)$, b^k units of commodity k must be shipped *from* node i . The destination node i of commodity k , noted $D(k)$, is called a *demand* node. At $D(k)$, b^k units of commodity k must be shipped *to* node i . When $b^k = 0$, node i is a *transshipment* node, which means that the amount of commodity k entering node i must always equal the amount of commodity k exiting node i . The decision variable x_{ij}^k is the amount of commodity k assigned to flow on arc (i, j) . Ahuja et al. [2] formulate the arc flow formulation as:

$$\text{MCNF-A} = \min \sum_{k \in K} \sum_{(i,j) \in A} c_{ij}^k x_{ij}^k \quad (3.1)$$

$$\text{s.t.} \quad \sum_{k \in K} x_{ij}^k \leq u_{ij} \quad \forall (i, j) \in A, \quad (3.2)$$

$$\sum_{\{j:(i,j) \in A\}} x_{ij}^k - \sum_{\{j:(j,i) \in A\}} x_{ji}^k = \begin{cases} b^k, & \text{if } i = O(k) \\ -b^k, & \text{if } i = D(k) \\ 0, & \text{otherwise} \end{cases} \quad \forall i \in N, k \in K, \quad (3.3)$$

$$x_{ij}^k \geq 0 \quad \forall (i, j) \in A, k \in K. \quad (3.4)$$

Constraints (3.2) restrict the total flow over an arc to be less than the arc capacity. Constraints (3.4) are referred to as *non-negativity constraints* and ensure that the amount of flow shipped on each arc is non-negative. *Balance of flow constraints* (3.3) ensure that the flow out of a node minus the flow into a node is equal to the supply/demand at the node. Balance of

flow constraints are equivalently written using matrix notation as $\mathcal{N}\mathbf{x}^k = \mathbf{b}^k$. The matrix \mathcal{N} is called the *node-arc incidence matrix*. Each column of \mathcal{N} corresponds to a single variable and contains the variable's constraint coefficients. Specifically, column \mathcal{N}_i will have a $+1$ in the i^{th} row and a -1 in the j^{th} row, indicating that the arc leaves node i and enters node j . All other entries in the column will be zero. The demand for commodity k at node i is contained in the i^{th} element of the vector \mathbf{b}^k , and the vector \mathbf{x}^k contains the decision variables.

3.5.1.2 Path Flow Formulation

A *path flow formulation* assigns flow to *paths* through the graph. Assigning flow to a path inherently assigns flow to all the arcs contained in the path. Thinking in terms of paths is often consistent with how transportation resources (e.g., aircraft, trucks, trains) are scheduled to operate in a real-world transportation system. For instance, when creating a channel route schedule, organic channel route schedulers assign an aircraft to a channel route mission (i.e., an entire path), rather than individual flight legs (i.e., an arc).

The following additional notation is required for the path flow formulation. Let P^k be the set of feasible paths for commodity k starting at the origin of k and terminating at the destination of k . The cost of path p is c_p^k , defined to be the sum of the costs of the individual arcs in the path: $c_p^k = \sum_{(i,j) \in A} \delta_{ij}^p c_{ij}^k$. The indicator δ_{ij}^p equals 1 if arc (i,j) is included in path p , and 0 otherwise. The decision variables x_p^k indicate the amount of commodity k assigned to flow on path p . Ahuja et al. [2] formulate the path flow formulation as:

$$\text{MCNF-P} = \min \sum_{k \in K} \sum_{p \in P^k} c_p^k x_p^k \quad (3.5)$$

$$\text{s.t.} \quad \sum_{k \in K} \sum_{p \in P^k} \delta_{ij}^p x_p^k \leq u_{ij} \quad \forall (i,j) \in A, \quad (3.6)$$

$$\sum_{p \in P^k} x_p^k = b^k \quad \forall k \in K, \quad (3.7)$$

$$x_p^k \geq 0 \quad \forall k \in K, p \in P^k. \quad (3.8)$$

Constraints (3.6) are the same as constraints (3.2) in MCNF-A. Constraints (3.7) ensure that the total flow of commodity k over all the paths for commodity k is equal to the demand of k . Alternately, constraints (3.7) can be written using *convexity constraints* by moving b^k to the left hand side of constraints (3.6) and changing the right hand side of constraints (3.7) to 1. In this reformulation, the decision variables x_p^k indicate the fraction of commodity k assigned to path p . Finally, constraints (3.8) ensure non-negative flows.




Figure 3-8: Network design choices

For graphs composed of many nodes, MCNF-P has fewer constraints than MCNF-A, but MCNF-A has fewer variables. MCNF-A has one variable for each commodity on every arc, a set of constraints for each arc, and a set of constraints for each node-commodity combination. MCNF-P only has a set of constraints for each arc and each commodity. A fewer number of constraints does not come without a price, however. MCNF-P has many more variables than MCNF-A because there are many feasible paths for each commodity.

3.5.2 Network Design Formulation

MCNF-A and MCNF-P determine the optimal flow of commodities through a graph with a fixed structure (i.e., the capacities of the arcs are fixed). However, there is a class of problems, called *network design problems*, in which the network structure is not fixed. In network design problems, *design* decisions are made about which arcs in the graph will exist to allow flow of the commodities. For example, in the channel route network, either a C-5 with a capacity of 50 tons or a C-141 with a capacity of 16 tons can fly between Dover AFB, DE (DOV) and Ramstein AB, Germany (RMS) (see *Figure 3-8*).

Additional notation is required for the network design formulation. Let F denote the set of resources that will provide the capacity on the arcs in the graph. The discrete *design variables* y_{ij}^f are associated with resources that may be allocated to the arcs in the graph. Each resource f provides capacity u_{ij}^f to arc (i, j) .

Often, there is a design cost d_{ij}^f associated with the use of one unit of resource f on arc (i, j) . Ahuja et al. [2], Magnanti and Wong [42], and Armacost [3] introduce network design formulations. We consider the following service network design formulation:

$$\text{NDF} = \min \sum_{k \in K} \sum_{p \in P^k} c_p^k x_p^k + \sum_{f \in F} \sum_{(i,j) \in A} d_{ij}^f y_{ij}^f \quad (3.9)$$

$$s.t. \quad \sum_{k \in K} \sum_{p \in P^k} \delta_{ij}^p b^k x_p^k \leq \sum_{f \in F} u_{ij}^f y_{ij}^f \quad \forall (i,j) \in A, \quad (3.10)$$

$$\sum_{p \in P^k} x_p^k = 1 \quad \forall k \in K, \quad (3.11)$$

$$\Phi y^f = 0 \quad \forall f \in F, \quad (3.12)$$

$$x_p^k \geq 0 \quad \forall k \in K, p \in P^k, \quad (3.13)$$

$$y_{ij}^f \in \{0,1\} \quad \forall (i,j) \in A, f \in F. \quad (3.14)$$

The objective function (3.9) minimizes the total commodity flow and resource flow costs. *Forcing constraints* (3.10) restrict the total flow on each arc to be less than the capacity assigned to the arc. *Convexity constraints* (3.11) ensure that all the demand is flown. Constraints (3.12) ensure balance of flow for the design variables. Balance of flow constraints are not generally found in network design problems, but are often introduced in transportation problems for which the design variables represent the flow of transportation resources such as aircraft, ships, or trucks. Non-negative cargo flows are enforced by constraints (3.13) and constraints (3.14) restrict the values of the design variables to be either zero or one. When we relax constraints (3.14) to allow an integral solution, the problem falls into the class of problems referred to as *network loading* (see Magnanti and Wong [42]).

3.6 Network Design Literature Review

The use of network design problems to model transportation systems is prevalent in the literature. Rather than providing an exhaustive literature review of network design problems, we point the reader to Magnanti and Wong [42], Gendron et al. [29], and Armacost [3]. Magnanti and Wong [42] provide an extensive overview of network design problems, with an emphasis on transportation problems. Their research includes optimization models, heuristic methods, and algorithms. Gendron et al. [29] provide a comprehensive survey of models and algorithms for network design problems, with an emphasis on telecommunications and transportation applications. Additionally, they develop several relaxation methods based on Lagrangean relaxation and differentiable optimization techniques (i.e., subgradient and bundle approaches). Armacost [3] provides an overview of network design problems, express shipment service network design, and solution strategies.

Rexing et al. [45] formulate a daily fleet assignment problem using time windows. They assign discretized time windows around the departure time of each flight arc and create copies of the flight arc within the surrounding time window. They formulate an optimization problem in which only one of the flight arcs in each time window can be chosen. They also present a series of network reduction techniques, based on the results from Hane et al. [34] to reduce the size of the network, and two algorithmic approaches for solving the model.

Grünert and Sabastien [33] identify planning tasks faced by postal and express shipment companies. They develop several optimization models to address these planning tasks, and discuss the use of the optimization models within a decision support system. Büdenbender et al. [17] develop a hybrid tabu search/branch-and-bound solution methodology to design a direct flight network for postal and express shipment companies.

Barnhart and Schneur [11] use a network design formulation, based solely on aircraft route variables, to design an air service network for express shipment service. Their research focuses on the design of pickup and delivery routes in a network consisting of a single hub. They also present a solution strategy based on column generation, and implement their model to evaluate various scenarios faced by a large commercial carrier.

Extending the single hub work of Barnhart and Schneur, Kim et al. [38] develop a multi-hub express shipment service network design formulation in terms of package flow variables and aircraft route variables. They develop various reduction techniques that reduce the model's size. They apply column generation techniques to deal with the large number of variables, and they develop a set of cutset inequalities to deal with the weak linear programming relaxations of the network design formulation. They solve the problem using a heuristic solution strategy.

3.7 Channel Route Formulations

This section describes two formulations we use to model the channel route network. Prior to describing the two formulations, we describe how commodities are classified. The two formulations are compared, which motivates a decomposition approach based on Dantzig-Wolfe decomposition.

3.7.1 Commodity Definition

There are many ways to classify the *cargo commodities* that flow through the channel route network. A cargo commodity is simply a collection of individual cargo items that have similar characteristics. The characteristics used to define a commodity are important because

the commodity definition must be consistent with the scope and fidelity of the real-world system, and consistent with the desired output. For instance, channel route planners assume all cargo can be palletized. Therefore, classifying cargo commodities by cargo category (i.e., bulk, oversize, outsize) or cargo type (e.g., rations, baggage) is not consistent with either the real-world system or the desired output.

In the channel route network, a commodity is defined by the following characteristics:

- ⇒ Origin aerial port
- ⇒ Destination aerial port
- ⇒ *Available to load time* (ALT). The ALT is the time that the cargo is available for pickup at the origin. It is assumed that the cargo is available at the beginning of the ALT.

All of the cargo items with the same origin, destination, and ALT are classified as a single cargo commodity. For example, cargo shipped from Dover AFB, DE (DOV) to Ramstein AB, Germany (RMS) on Tuesday is different than cargo shipped from DOV to RMS on Saturday. Each cargo commodity has an associated tonnage, which is the sum of the weights of all the individual cargo items in the cargo commodity. For instance, consider ten tons of rations that must be shipped from DOV to RMS on Tuesday and twenty-five tons of baggage that must be shipped from DOV to RMS on Tuesday. All of these cargo items are classified as one cargo commodity: thirty-five tons from DOV to RMS on Tuesday.

3.7.2 Channel Route Arc-Flow Formulation

This section introduces the first formulation used to model the channel route network. The *arc flow channel route model* is concerned with assigning cargo and aircraft to individual arcs in the graph. The optimization solver must determine the optimal set of arcs to use in order to flow each cargo commodity through the graph. In order to flow cargo over an arc, however, an aircraft with sufficient capacity must also be scheduled to traverse the arc. The objective of the formulation is to move all the cargo commodities using the fewest number of aircraft, subject to constraints placed on the cargo flows and the aircraft flows. Although we introduce the formulation with an objective of minimizing aircraft, we evaluate various objective functions in later chapters.

The cargo flows and aircraft flows are represented with different variables in the formulation. Cargo flows and aircraft flows must be represented with different variables because cargo and aircraft flow through the system following different paths and are subject to different

sets of constraints. Each cargo flow variable corresponds to the flow of one commodity type on a single arc in the graph. The aircraft flow variables are decomposed into aircraft flows over flight arcs and aircraft flows over ground arcs. The aircraft flow variables are separated because aircraft traversing ground arcs are subject to different constraints than aircraft traveling over flight arcs. Each aircraft flow variable corresponds to the flow of one aircraft type over a single ground arc or flight arc in the graph.

The following notation is used in the arc flow formulation:

SETS

PHYSICAL ASSETS IN THE CHANNEL ROUTE NETWORK (AERIAL PORTS & AIRCRAFT)

B^S	set of all aerial ports $i \in B^S$ that have an assigned aircraft squadron;
B^{NS}	set of all aerial ports $i \in B^{NS}$ that do not have an assigned aircraft squadron;
B	set of all aerial ports $i \in B, j \in B, (B = B^S \cup B^{NS})$;
M	set of all aircraft models $m \in M$ (e.g., Boeing C-5 Galaxy);
F	set of all aircraft types $f \in F$. Each aircraft type f is indexed by $\{m, i\}$, indicating the model type $m \in M$ and its squadron aerial port $i \in B^S$.
F^m	set of all aircraft types $f \in F$ with model type $m \in M$; and
T_i^{ops}	set of all time periods $t \in T_i^{ops}$ in which aerial port $i \in B$ is closed as a result of operations hours, BASH hours, and/or quiet hours.

NETWORK STRUCTURE

A	set of all arcs $a \in A$. Each arc a is indexed by $\{i, j, t\}$, indicating that a originates at aerial port $i \in B$ in time $t \in T$, and terminates at aerial port $j \in B$. Let $O(a)$ be the origin aerial port $i \in B$ of arc $a \in A$ and $D(a)$ be the destination aerial port $j \in B$ of arc $a \in A$;
\bar{A}	set of flight arcs $a \in \bar{A} (\bar{A} \subset A)$;
\underline{A}	set of ground arcs $a \in \underline{A} (\underline{A} \subset A)$;
N	set of all nodes $n \in N$. Each node n is indexed by $\{i, t\}$ indicating the aerial port $i \in B$ and time period $t \in T$. Let $B(n)$ be the aerial port $i \in B$ associated with node $n \in N$;
N^{B^S}	set of squadron base nodes $n \in N^{B^S}, B(n) \in B^S$;
$N^{B^{NS}}$	set of squadron base nodes $n \in N^{B^{NS}}, B(n) \in B^{NS}$;
$A[n^+]$	set of all arcs incoming to node $n \in N$;
$A[n^-]$	set of all arcs outgoing from node $n \in N$;

$A[t^c]$ set of all arcs that cross the cutset time t^c . $A[t^c]$ is called the *cutset*; and
 T set of all time periods $t \in T$.

COMMODITIES

K set of all cargo commodities $k \in K$. Each commodity k is indexed by $\{i, j, t\}$, indicating that cargo must move from aerial port $i \in B$ to aerial port $j \in B$ and that the cargo is available to load at aerial port $i \in B$ in time period $t \in T$.

DATA

u^f capacity of aircraft type $f \in F$;
 u_a^{FB} capacity of fixed-buy commercial aircraft assigned to arc $a \in A$;
 b^k demand for cargo commodity $k \in K$;
 ρ^f number of available aircraft type $f \in F$;
 H^m required flying hours for aircraft model $m \in M$;
 \bar{H}^m upper limit on the number of required flying hours for aircraft model $m \in M$;
 \underline{H}^m lower limit on the number of required flying hours for aircraft model $m \in M$;
 h_a^f number of required flying hours counting towards H^m when aircraft type $f \in F$ is assigned to arc $a \in A$;
 t^c cutset time period;
 e_n maximum number of aircraft that can be serviced simultaneously at node $n \in N$;
 \mathcal{D} node-arc incidence matrix for aircraft flow variables; and
 \mathcal{N} node-arc incidence matrix for cargo flow variables.

DECISION VARIABLES

x_a^k amount of commodity $k \in K$ to send on arc $a \in A$;
 y_a^f number of aircraft type $f \in F$ assigned to flight arc $a \in \bar{A}$; and
 g_a^f number of aircraft type $f \in F$ assigned to ground arc $a \in \underline{A}$.

A Note on Notation

For convenience, shorthand notation is used for sets that are an intersection of two or more sets. For example, $\bar{A} \cap A[n^+]$ yields the set of all incoming flight arcs into node n and is

written as $\bar{A}[n^+]$. Similarly, $\bar{A}[t^c]$ is the set of flight arcs in cutset $A[t^c]$, and is a result of the intersection $\bar{A} \cap A[t^c]$.

FORMULATION

$$\text{CRM-A} = \min \sum_{f \in F} \sum_{a \in \bar{A}[t^c]} y_a^f + \sum_{f \in F} \sum_{a \in \underline{A}[t^c]} g_a^f \quad (3.15)$$

$$\text{s.t.} \quad \sum_{k \in K} x_a^k \leq u_a^{FB} + \sum_{f \in F} u^f y_a^f \quad \forall a \in \bar{A}, \quad (3.16)$$

$$\mathbf{Kx}^k = \mathbf{b}^k \quad \forall n \in N, k \in K, \quad (3.17)$$

$$\sum_{a \in \underline{A}[t^c]} g_a^f + \sum_{a \in \bar{A}[t^c]} y_a^f \leq \rho^f \quad \forall f \in F, \quad (3.18)$$

$$\Phi \mathbf{y}^f = \mathbf{0} \quad \forall f \in F, \quad (3.19)$$

$$\sum_{f \in F^m} \sum_{a \in \bar{A}} y_a^f h_a^f \leq \bar{H}^m \quad \forall m \in M, \quad (3.20)$$

$$\sum_{f \in F^m} \sum_{a \in \bar{A}} y_a^f h_a^f \geq \underline{H}^m \quad \forall m \in M, \quad (3.21)$$

$$\sum_{f \in F} \sum_{a \in \bar{A}[n^+]} y_a^f \leq e_n \quad \forall n \in N^{B^s}, \quad (3.22)$$

$$\sum_{f \in F} g_a^f \leq e_n \quad \forall n \in N^{B^{ns}}, a \in \underline{A} : O(a) = B(n) \quad (3.23)$$

$$y_a^f = 0 \quad \forall a \in \bar{A} : t \in T^{ops}, f \in F, \quad (3.24)$$

$$x_a^k \geq 0 \quad \forall k \in K, a \in A, \quad (3.25)$$

$$y_a^f, g_a^f \in \mathbb{Z}^+ \quad \forall f \in F, a \in A. \quad (3.26)$$

Objective (3.15) minimizes the number of aircraft assigned to either a ground or a flight arc in the cutset. *Forcing constraints* (3.16) restrict the amount of cargo that can flow over a flight arc to be less than or equal to the capacity of the aircraft traversing the arc. The aircraft capacity on an arc is determined by the military aircraft assigned to the arc and the commercial aircraft scheduled to traverse the arc. The commercial aircraft capacity on each arc is not a design decision and is known a priori based on the set of fixed-buy missions. Furthermore, the amount of cargo is only restricted on the flight arcs, which captures the assumption that aerial ports have unlimited cargo capacity. Constraints (3.17) and (3.19) conserve the flow of cargo and aircraft, respectively.

The *aircraft count constraints* (3.18) restrict the number of aircraft assigned to the ground and flight arcs that cross the cutset time to be less than the number of aircraft available. Constraints (3.20) and (3.21) ensure that the required number of flying hours are flown for each aircraft model. Although the annual number of flying hours is known exactly, the number of flying hours required per month for each aircraft model is an estimate. Thus, constraints (3.20)

and (3.21) ensure that the number of flying hours flown is within an acceptable window of the estimated number of flying hours that must be flown during the planning horizon.

Constraints (3.22) and (3.23) enforce the working maximum on ground (WMOG) limitations at each aerial port. At aerial ports that have an assigned aircraft squadron, it is necessary to enforce these constraints on the flight arcs so that the total number of available aircraft stationed at the aerial port is not restricted to the size of WMOG. At aerial ports that do not have an assigned aircraft squadron, the constraint is enforced on the ground arc.

Constraints (3.24) enforce the operating hours, bird air strike hours (BASH), and quiet hours at each aerial port. These constraints are enforced on the flight arcs, capturing the fact that aircraft cannot fly into, or out of, the aerial port during these time periods. These constraints are not enforced on the ground arcs, allowing the aircraft located at the aerial port prior to these time periods to remain on the ground while the aerial port is closed. Finally, cargo flows must be non-negative (3.25) and the number of aircraft assigned must be integer (3.26).

DISCUSSION OF CRM-A

CRM-A has three weaknesses. The first weakness is the size of CRM-A for realistic problem instances. The second is CRM-A's inability to model frequency channels and complex operational rules and regulations governing aircraft and aircrews. Finally, the third weakness of CRM-A is its weak linear programming (LP) relaxation. The following sections discuss each of these weaknesses in detail.

Formulation Size

CRM-A is tractable for small problem instances. However, the size of CRM-A becomes extremely large with only small changes in the problem inputs. For instance, consider a channel route network with the characteristics in *Table 3-1*.

The size of CRM-A is sensitive to changes in the number of days in the planning horizon, the length of each period in the planning horizon, and the number of cargo commodities (see *Table 3-2*). Computational experience on similarly sized network design problems with the same structure (see Kim, et al. [38]) demonstrate that CRM-A would be intractable for a thirty-day planning horizon.

Squadron Aerial Ports	5
Non-Squadron Aerial Ports	25
Aircraft Types	5

Table 3-1: Example Channel Route Network Characteristics

	COMMODITIES	PLANNING HORIZON (DAYS)	NUMBER OF CONSTRAINTS	INTEGER VARIABLES	TOTAL VARIABLES
2 periods per day (1 period = 12 hrs)	5	3	6,891	26,100	52,200
	8	4	9,893	34,800	52,200
	20	7	22,319	60,900	304,500
	254	30	516,705	261,000	13,519,800
6 periods per day (1 period = 4 hrs)	5	3	20,583	78,300	156,600
	8	4	29,589	104,400	271,440
	20	7	66,867	182,700	913,500
	254	30	1,550,025	783,000	40,559,400

Table 3-2: Size of CRM-A for various Channel Route Network instances

Operational Rules

Another weakness of CRM-A is its inability to model frequency channels and many of the operational rules and regulations that govern aircraft and aircrews in the real-world channel route network. Representing the operational rules and regulations as algebraic constraints is cumbersome. For instance, aircrews typically remain with an aircraft for the duration of an entire channel route mission. As a result, the flow of aircraft through the graph should be constrained by the aircrew limitations presented in *Table 2-3*. However, enforcing these constraints on an individual arc is not possible because the rules apply to a collection of arcs. A limited number of the operational rules and regulations can be built into the structure of the graph, using the process described in §3.4.2. For example, constraints (3.24) can be removed from the formulation and enforced by excluding these arcs during the construction of the graph.

LP-Relaxation

The design variables in CRM-A are constrained to be integer by the integrality constraints (3.26). To obtain an integer solution, the integrality constraints are relaxed and the resulting *linear programming relaxation* (LP-relaxation) is solved. If the design variables in the LP-relaxation solution are fractional, then a branch-and-bound algorithm is solved to obtain integer values for the design variables. The ability to successfully solve large MIP problems hinges on the quality of the LP-relaxation. The quality of the LP-relaxation is determined by how well it approximates the convex hull of the MIP [8].

Network design formulations that have the same structure as CRM-A typically have weak LP-relaxations because of the forcing constraints (3.16) and the aircraft balance of flow constraints (3.19). The objective of CRM-A is to minimize the number of aircraft used.

Consequently, the LP-relaxation of CRM-A has no incentive to assign an entire aircraft to flow on an arc if the cargo traversing that arc is less than the aircraft's capacity. This causes fractional aircraft to be assigned in the LP-relaxation. The problem of fractional aircraft is exacerbated by the aircraft balance of flow constraints because fractional aircraft are propagated throughout the network. Ultimately, this leads to a LP-relaxation that has many fractional aircraft flow variables. A highly fractional LP-relaxation provides a poor lower bound on the optimal integer solution, which makes the integer program (IP) difficult and more time-consuming to solve.

3.7.3 Channel Route Path Flow Formulation

This section presents an alternative formulation to the arc flow formulation. The *path flow channel route model* assigns cargo and aircraft to *paths* through the graph. A *path* is simply a collection of arcs in the graph that represents a feasible sequence of movements for cargo or aircraft. Thus, rather than making a distinct decision about the flow on each arc (i.e., the arc formulation), a decision to use a path simultaneously makes the decision to use all the arcs contained in the path. In the following formulation the paths associated with cargo and aircraft are referred to as *cargo paths* and *aircraft routes*, respectively.

The following additional notation is required.

SETS

COMMODITIES

W set of all frequency channel commodities $w \in W$. Each frequency channel w is indexed by $\{i, j, t, t'\}$, indicating the origin aerial port $i \in B$, the destination aerial port $j \in B$, and the time period – starting at $t \in T$ and ending at $t' \in T$ – during which the requirement must be satisfied.

CARGO PATHS

P set of all cargo paths $p \in P$; and
 P^k set of all cargo paths for commodity k , ($P^k \subset P$).

AIRCRAFT ROUTES

R set of all aircraft routes $r \in R$; and
 R^f set of all aircraft routes $r \in R$ that can be flown by aircraft type $f \in F$ ($R^f \subset R$).

DATA

- ψ_w^{FB} number of times that frequency channel $w \in W$ is covered by a commercial fixed-buy mission; and
- ξ^w number of times frequency requirement $w \in W$ must be covered.

INDICATOR VARIABLES

$$\delta_a^p = \begin{cases} 1, & \text{if cargo path } p \in P \text{ contains arc } a \in A \\ 0, & \text{otherwise;} \end{cases}$$

$$\delta_a^r = \begin{cases} 1, & \text{if aircraft route } r \in R \text{ contains arc } a \in A \\ 0, & \text{otherwise;} \text{ and} \end{cases}$$

$$\delta_w^r = \begin{cases} 1, & \text{if aircraft route } r \in R \text{ covers frequency requirement } w \in W \\ 0, & \text{otherwise.} \end{cases}$$

DECISION VARIABLES

- x_p^k fraction of type $k \in K$ cargo to send on path $p \in P$; and
- y_r^f number of aircraft type $f \in F$ assigned to fly on route $r \in R$.

FORMULATION

$$\text{CRM-P} = \min \sum_{f \in F} \sum_{r \in R^f} \sum_{a \in A[k^c]} \delta_a^r y_r^f \quad (3.27)$$

$$s.t. \quad \sum_{k \in K} \sum_{p \in P^k} \delta_a^p b^k x_p^k \leq u_a^{FB} + \sum_{f \in F} \sum_{r \in R^f} \delta_a^r u^f y_r^f \quad \forall a \in \bar{A}, \quad (3.28)$$

$$\sum_{p \in P^k} x_p^k = 1 \quad \forall k \in K, \quad (3.29)$$

$$\sum_{r \in R^f} \sum_{a \in A[k^c]} \delta_a^r y_r^f \leq \rho^f \quad \forall f \in F, \quad (3.30)$$

$$\Phi y^f = 0 \quad \forall f \in F, \quad (3.31)$$

$$\sum_{f \in F^m} \sum_{r \in R^f} \sum_{a \in \bar{A}} \delta_a^r h_a^f y_r^f \leq \bar{H}^m \quad \forall m \in M, \quad (3.32)$$

$$\sum_{f \in F^m} \sum_{r \in R^f} \sum_{a \in \bar{A}} \delta_a^r h_a^f y_r^f \geq \underline{H}^m \quad \forall m \in M, \quad (3.33)$$

$$\sum_{f \in F} \sum_{r \in R^f} \sum_{a \in \bar{A}} \delta_a^r y_r^f \leq e_n \quad \forall n \in N^{B^s}, \quad (3.34)$$

$$\sum_{f \in F} \sum_{r \in R^f} \sum_{a \in \bar{A}} \delta_a^r y_r^f \leq e_n \quad \forall n \in N^{B^{ns}}, \quad (3.35)$$

$$\sum_{f \in F} \sum_{r \in R^f} \delta_w^r y_r^f + \psi_w^{FB} \geq \xi^w \quad \forall w \in W, \quad (3.36)$$

$$x_p^k \geq 0 \quad \forall k \in K, p \in P^k, \quad (3.37)$$

$$y_r^f \in \mathbb{Z}^+ \quad \forall f \in F, r \in R^f. \quad (3.38)$$

Equation (3.27) minimizes the number of aircraft assigned to the arcs in the cutset. Forcing constraints (3.28) ensure that the amount of cargo assigned to an arc does not exceed the amount of commercial aircraft capacity on that arc plus the amount of military aircraft capacity assigned to that arc. Convexity constraints (3.29) ensure that all the demand for each cargo commodity is shipped from its origin to its destination. Constraints (3.30) are the aircraft count constraints. Constraints (3.31) are the balance-of-flow constraints for aircraft. Constraints (3.32) and (3.33) ensure that the required flying hours are flown for each aircraft model. Constraints (3.34) and (3.35) enforce the WMOG restriction at each base. Constraints (3.36) ensure that the frequency channels are serviced by either a military aircraft or a commercial aircraft. Finally, cargo flows must be non-negative (3.37) and the number of aircraft assigned must be integer (3.38).

VARIABLES AS COLUMNS

The cargo path and aircraft route variables are represented by columns in CRM-P. This is illustrated in *Example 3-3*.

Example 3-3: Consider the aircraft route (flown by an aircraft with 50 tons of capacity) and the cargo paths in Figure 3-9, and the following requirements and frequency channels:

CHANNEL TYPE	ORIGIN	DESTINATION	ALT/FREQ	TONS
Requirements	A	D	0	20
Requirements	C	B	1	15
Frequency	A	B	1/week	--

The aircraft route and cargo path variables are represented as columns in CRM-P by placing the appropriate non-zero coefficients in the rows that correspond to the constraints in which the variable appears in the formulation. The remainder of the entries in the column are zero. The aircraft route and cargo paths columns are represented as illustrated in Figure 3-10.




Figure 3-9: Aircraft Routes and Cargo Paths as columns in CRM-P

DISCUSSION OF CRM-P

CRM-P relies on the ability to identify feasible cargo paths and aircraft routes. Assume that there is an oracle that generates cargo paths and aircraft routes offline, and submits them to CRM-P as an input. An advantage of using path variables is that the oracle can enforce the difficult operational rules and regulations that CRM-A lacked the ability to enforce, which alleviates the need to explicitly include these constraints in CRM-P. Furthermore, using path variables allows the frequency channels to be modeled.

Cargo Path 1	Cargo Path 2	Aircraft Route			
20	0	-50	\leq	0	$(A,0) \rightarrow (C,1)$
20	15	-50	\leq	0	$(C,1) \rightarrow (D,2)$
0	15	-50	\leq	0	$(D,2) \rightarrow (B,3)$
0	0	-50	\leq	0	$(B,3) \rightarrow (A,4)$
1	0	0	$=$	1	$A \rightarrow D$
0	1	0	$=$	1	$C \rightarrow B$
0	0	1	\geq	1	$A \rightarrow B$

Flight Arcs

Commodities

Freq Channel

Figure 3-10: Structure of columns in CRM-C

	Instance 1	Instance 2	Instance 3	Instance 4
<i>Commodities</i>	5	8	20	254
<i>Time Horizon (days)</i>	3	4	7	30
<i>Constraints in CRM-P</i>	6,063	8,048	14,006	59,778
<i>Constraints in CRM-A</i>	6,891	9,893	22,319	516,705

Table 3-3: Number of constraints in CRM-P for various Channel Route Network instances

For large problem instances, CRM-P contains significantly fewer constraints than CRM-A. Table 3-3 compares the number of constraints in CRM-P and CRM-A for the channel route network described in Table 3-1 and with two time periods per day.

The drawback to using path variables is that there are many feasible cargo paths and aircraft routes, which results in an extremely large number of columns. The large number of columns can be addressed, however, using column generation. Column generation is based on the observation that only a fraction of the total number of variables in a math programming formulation will be non-zero at an optimal solution. Therefore, the key idea of column generation is to solve the problem over a limited number of the total columns. If the solution is optimal when solved over the limited set of columns, there is no need to consider the columns that are not in the problem formulation. However, if the solution is not optimal, a cost-improving column (i.e., a column with a favorable reduced cost) is identified and added to the problem formulation. The problem is resolved over the slightly larger set of columns. This process continues until an optimal solution is reached. The success of column generation relies on the ability to identify cost-improving columns without explicitly examining every feasible column [2]. This is accomplished by taking advantage of the structural properties of a column, which typically involves solving another easier optimization problem.

3.8 Decomposition

The purpose of this section is to develop a decomposition approach aimed at improving the tractability of CRM-P (i.e., addressing the large number of variables). The decomposition is based on classical Dantzig-Wolfe decomposition using column generation ([2],[15]). The main advantage of Dantzig-Wolfe decomposition is computational. Namely, it allows a single, large problem to be decoupled into several smaller, more tractable problems. In addition to the computational benefits, Dantzig-Wolfe decomposition is also a natural representation of decentralized planning organizations, such as the TACC.

In Dantzig-Wolfe decomposition, a *master problem* coordinates the use of shared resources among smaller *subproblems* so that the overall distribution of resources among the subproblems is optimal for the entire system. In the channel route planning problem, sub-planners, called the *cargo path generator* and the *aircraft route generator*, identify cargo paths for each cargo commodity and aircraft routes for each aircraft type to be considered by the *master level scheduler*. The master level scheduler determines the best subset of the proposed cargo paths and aircraft routes, subject to the constraints placed on the cargo commodities and the aircraft. Additionally, the master level scheduler determines the value (i.e., dual price) associated with each of the shared resources. The sub-planners (i.e., sub-problems), in turn, use the dual prices of the shared resources to identify additional cost improving cargo paths and aircraft routes. The master level scheduler and the sub-problems iterate in this fashion until the system-wide optimal set of cargo paths and aircraft routes has been identified (see *Figure 3-11*).

When applied to MCNF problems, there exists one subproblem for each commodity type (see Ahuja et al. [2]). The channel route problem contains a MCNF problem, so there exists one subproblem for each cargo commodity. All commodity subproblems are grouped together and referred to as the *cargo path generator*. In order for a cargo commodity to use an arc in the graph, an aircraft with sufficient capacity must be assigned to that arc. Therefore, a subproblem is also solved for each aircraft type. All of the aircraft subproblems are grouped together and referred to as the *aircraft route generator*. Clark [20] uses this type of decomposition to solve a large-scale munitions logistics problem in which the master problem determines the optimal mix of ammunition paths and ship paths through a time-space graph spanning a 30-day planning horizon.

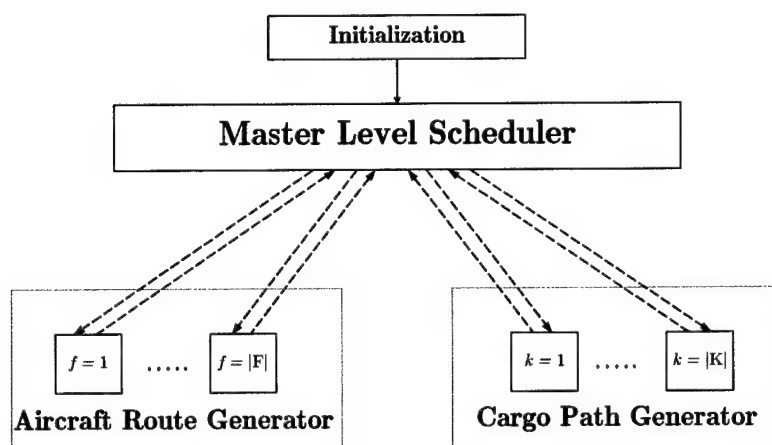


Figure 3-11: Dantzig-Wolfe Decomposition applied to CRM-P

Not only does the decomposition in *Figure 3-11* divide the channel route planning problem into several smaller, more tractable problems, it also resembles the decentralized nature of the current channel route planning process. In the framework of the decomposition, the aircraft route generator performs the role of the organic channel scheduler, and the cargo path generator plays the role of the CONUS cargo bookie. Although the organic channel scheduler and the CONUS cargo bookie iterate during the current channel route planning process, they do so in a sequential manner without utilizing a master level scheduler that has a global view of the system. The remainder of this section develops the mathematical formulations for the master level scheduler, the aircraft route generator, and the cargo path generator.

3.8.1 Relaxed Constraints

To simplify the decomposition, we deal with a special case of CRM-P. The following simplifying assumptions are made in order to make an initial assessment of this decomposition's viability for modeling the real-world channel route network. These assumptions will be relaxed in later chapters.

Assumption 1: The planning horizon is no more than one week

Limiting the planning horizon leads to the following assumptions, which allow several sets of constraints in CRM-P to be relaxed.

Assumption 2: Unlimited Aircraft, $\rho^f = \infty$

The contract number of aircraft provided to the Barrelmaster serves as an upper limit on the number of aircraft that can be used to fly channel route missions during the execution month. The contract number of aircraft could also be used as an upper limit for a weekly problem; however, a channel route schedule that uses all of the available aircraft in one week is not likely. Therefore, it is assumed that a weekly channel route schedule will not yield a solution in which the aircraft count constraints are binding. As a result, the aircraft count constraints (3.30) are relaxed.

Assumption 3: Unrestricted Flying Hours, $\underline{H}^m = 0$ and $\overline{H}^m = \infty$

The annual flying hours program requirements are disaggregated into monthly target level estimates. AMC uses the monthly target level estimates to gauge if each of the aircraft models will meet its annual flying hours requirement.

However, weekly target level estimates do not exist. As a result, flying hour constraints (3.32) and (3.33) are relaxed.

Assumption 4: Unlimited WMOG capacity, $e_n = \infty$

Assuming that each aerial port has unlimited WMOG capacity relaxes constraints (3.34) and (3.35). Channel route planners currently ignore WMOG capacity when creating the initial cut and assume that WMOG conflicts can be resolved within several days of the start of each mission.

Assumption 5: Cargo Paths and Aircraft Routes satisfy balance of flow

It is assumed that the cargo path generator and the aircraft route generator identify cargo paths and aircraft routes that span the entire planning horizon. Furthermore, we assume that by construction, the cargo paths and aircraft routes satisfy balance of flow at each node in the graph. As a result, the balance of flow constraints (3.31) are relaxed.

3.8.2 Master Level Scheduler

Based on the assumptions in the previous section and using the notation already introduced, the master level scheduler is written as:

$$\text{CRS-ML} = \min \sum_{f \in F} \sum_{r \in R^f} \sum_{a \in A[t^c]} \delta_a^r y_r^f \quad (3.39)$$

$$s.t. \sum_{k \in K} \sum_{p \in P^k} \delta_a^p b^k x_p^k - u_a^{FB} - \sum_{f \in F} \sum_{r \in R^f} \delta_a^r u^f y_r^f \leq 0 \quad \forall a \in \bar{A} \quad [\pi_a^{\bar{A}}], \quad (3.40)$$

$$\sum_{k \in K} \sum_{p \in P^k} \delta_a^p b^k x_p^k + u_a^{FB} + \sum_{f \in F} \sum_{r \in R^f} \delta_a^r u^f y_r^f \geq 0 \quad \forall a \in \underline{A} \quad [\pi_a^{\underline{A}}], \quad (3.41)$$

$$\sum_{p \in P^k} x_p^k = 1 \quad \forall k \in K \quad [\sigma_k], \quad (3.42)$$

$$\sum_{f \in F} \sum_{r \in R^f} \delta_w^r y_r^f + \psi_w^{FB} \geq \xi^w \quad \forall w \in W \quad [\alpha_w], \quad (3.43)$$

$$x_p^k \geq 0 \quad \forall k \in K, p \in P^k, \quad (3.44)$$

$$y_r^f \in \mathbb{Z}^+ \quad \forall f \in F, r \in R^f. \quad (3.45)$$

Constraints (3.40) and constraints (3.42) through (3.45) are the same as the constraints in CRM-P. Constraints (3.41) are added to aid the development of the subproblems in the next

section. These constraints capture the fact that the amount of flow on a ground arc is not restricted by the aircraft capacity assigned to that arc. The dual prices associated with each set of constraints are indicated in brackets.

3.8.3 Shortest Path Subproblems

Given dual prices from CRS-ML, the role of the subproblems is to identify cost improving columns to be considered by CRS-ML. We determine if such a column exists by solving a shortest path problem for each cargo commodity and each aircraft type. The objective of each subproblem is to find the cargo path (aircraft route) with the least reduced cost. If the reduced cost is negative, it is a cost improving column and it is added to CRS-ML. The success of the decomposition relies on the ability to solve each subproblem quickly by taking advantage of the underlying structure of the graph. Specifically, the graph associated with each subproblem is acyclic, which means that the shortest path can be determined quickly using a topological search algorithm (see Ahuja et al. [2]).

3.8.3.1 Cargo Path Generator

Each cargo commodity subproblem utilizes dual price information from CRS-ML to generate a new cargo path. A new cargo path is added to CRS-ML if its reduced cost is negative. The reduced cost of a cargo path column is:

$$\bar{c}_{x_p^k} = \sum_{a \in \bar{A}} \delta_a^p \pi_a^{\bar{A}} - \sum_{a \in \underline{A}} \delta_a^p \pi_a^{\underline{A}} - \sigma_k. \quad (3.46)$$

The objective of cargo commodity subproblem k is to find a cargo path $p \in P^k$ through a modified graph that minimizes the objective function:

$$\min z^k = \sum_{a \in \bar{A}} \delta_a^p \pi_a^{\bar{A}} - \sum_{a \in \underline{A}} \delta_a^p \pi_a^{\underline{A}} - \sigma_k. \quad (3.47)$$

The constant σ_k can be removed from the objective without influencing the values of the decision variables at the optimal solution. If $(z^k - \sigma^k) < 0$, the column has negative reduced cost and is added to the master level scheduler.

A new graph, which is a modified version of the graph presented in *Figure 3-7f*, is created for each subproblem. The portion of the graph considered in each subproblem is restricted to ensure that the shortest path algorithm begins at the APOE of the commodity

during the proper available to load time, and terminates at the APOD of the commodity. A feasible path through the modified graph is also a feasible path through the original graph.

For each arc in the modified graph there is one decision variable δ_a^p that indicates if arc a is included in the cargo path p . Each flight arc is assigned a modified cost $\pi_a^{\bar{A}}$. Likewise, each ground arc is assigned a modified cost $\pi_a^{\underline{A}}$. The constraints (3.41) will always be non-binding, so the dual cost associated with each ground arc will always be zero.

To solve for the shortest path, the modified graph is topologically ordered and a standard reaching or pulling algorithm (see Ahuja et al. [2]) is applied. Both the reaching and pulling algorithms visit each arc in the graph one time and will, therefore, solve the shortest path in $O(A)$ time, where A is the set of arcs in the graph.

3.8.3.2 Aircraft Route Generator

Similar to the cargo commodity subproblems, a shortest path subproblem is solved for each aircraft type. Each subproblem uses dual prices from CRS-ML to identify if a negative reduced cost aircraft route column exists. The reduced cost of an aircraft route column is:

$$\bar{c}_{y_r^f} = 1 - u^f \left(\sum_{a \in \bar{A}} \delta_a^r \pi_a^{\bar{A}} + \sum_{a \in \underline{A}} \delta_a^r \pi_a^{\underline{A}} \right) - \sum_{w \in W} \delta_r^w \alpha_w. \quad (3.48)$$

The objective of aircraft subproblem f is to find an aircraft route $r \in R^f$ through a modified graph that minimizes the objective function:

$$\min z^f = 1 - u^f \left(\sum_{a \in \bar{A}} \delta_a^r \pi_a^{\bar{A}} + \sum_{a \in \underline{A}} \delta_a^r \pi_a^{\underline{A}} \right) - \sum_{w \in W} \delta_r^w \alpha_w, \text{ which is equivalent to:} \quad (3.49)$$

$$1 - \left(\min z^f = u^f \left(\sum_{a \in \bar{A}} \delta_a^r \pi_a^{\bar{A}} + \sum_{a \in \underline{A}} \delta_a^r \pi_a^{\underline{A}} \right) + \sum_{w \in W} \delta_r^w \alpha_w \right). \quad (3.50)$$

If $(1 - z^f) < 0$, then the aircraft route column has negative reduced cost and it is added to the master level scheduler.

A modified graph is constructed for each aircraft type. The modified graph is restricted to ensure that the shortest path algorithm begins at the aircraft's squadron aerial port, at the beginning of the planning horizon, and terminates at the aircraft's squadron aerial port by the end of the planning horizon. The flight arcs are assigned a modified cost $\pi_a^{\bar{A}}$, and the ground

arcs are assigned a modified cost π_a^A . The dual prices of the ground arcs will always be zero because constraints (3.41) will always be non-binding. The decision variables δ_a^r indicate if arc a is included in aircraft route r . The decision variables δ_r^w indicate if frequency requirement w covered by aircraft route r .

The decision variables δ_r^w complicate the shortest path algorithm because it is not possible to assign these decision variables to individual arcs in the graph. As a result, the modified frequency channel costs α_w cannot simply be distributed across individual arcs like the modified arc costs. Therefore, a multilabel shortest path algorithm must be used (see Barnhart et al. [9] and Desrochers [26]). Multilabel shortest path problems require that information about all *non-dominated* paths be kept at each node in the graph. Given two paths, a path is dominated when there exists a less expensive path and one of paths does not have the potential to cover one or more frequency requirements that the other path does not have the potential to cover (and vice versa).

In general, multilabel shortest paths problems take longer to solve than pure shortest path problems. The time required to solve multilabel shortest path problems is a function of the number of labels that must be maintained at each node in the graph, and the number of dominance rules that must be enforced at each node in the graph. More labels and dominance rules result in longer solve times. There are relatively few labels in the aircraft route subproblems because there are relatively few frequency channels. Additionally, only two dominance rules must be enforced. Overall, the multilabel shortest path algorithm does not require a prohibitive amount of time to solve, which means that aircraft route columns can be identified quickly. In fact, the modified graph is topologically ordered and each arc is visited only once during the multilabel shortest path algorithm. The maximum number of labels at any node is equal to the total number of frequency requirements (W). The number of arcs (A) in the graph is always much larger than the number of nodes (N) in the graph. The running time of the multilabel shortest path algorithm is $O(A+NW)$. However, because A dominates NW for this problem, the running time of the multilabel shortest path is $O(A)$.

MULTILABEL SHORTEST PATH ALGORITHM FOR THE AIRCRAFT ROUTE SUBPROBLEMS

A standard multilabel shortest path algorithm (similar to the one used by Barnhart et al. [9] to solve the crew pairing subproblems of an airline crew scheduling formulation) is used to solve the aircraft route subproblems. The multilabel shortest path algorithm requires that information about the frequency requirements be kept at each node. Each label l at node n consists of the following information:

- $\Rightarrow p$: predecessor node;
- $\Rightarrow c(n, l)$: cost of the shortest path to node n , label l . The cost equation is:

$$c(n, l) = c^{l(p)} + \pi_a + \delta_n^w \alpha_w, \text{ where} \quad (3.51)$$

$c^{l(p)}$ = the cost of the shortest path at predecessor node label $l(p)$;

π_a = dual price of arc a between the predecessor node p and the current node n ;

α_w = dual price of frequency requirement w ; and

$$\delta_n^w = \begin{cases} 1, & \text{if path to node } n \text{ covers frequency requirement } w \in W \\ 0, & \text{otherwise;} \end{cases}$$

- $\Rightarrow L^o$: a list of the frequency channel origin aerial ports visited in the shortest path to node n ;
- $\Rightarrow l(p)$: predecessor node label number that the current label corresponds to; and
- $\Rightarrow l$: current label number.

The labels at each node are updated according the procedure in *Figure 3-12*. It should be noted that the multilabel shortest path problem can be solved quickly because the number of frequency requirements is relatively small. Specifically, $|W| = 19$ for the time horizon considered as a result of Assumption 1 stated above.

There are two rules used to determine if a label dominates another label. Rule 1 applies only to labels at nodes in the graph that are not in the last time period. Rule 2 applies only to labels at nodes in the last time period.

RULE 1: Consider two labels. If L^o is empty for both labels, then the label with the greatest cost is dominated and can be discarded.

RULE 2: Consider two or more labels at a node in the last time period of the graph. The label with the least cost dominates all other labels, which can be discarded.

An example of how the labels are updated and how the dominance rules are applied is illustrated in *Figure 3-13*. For simplicity, the graph only considers a select number of arcs.


```

W ← List of frequency requirements
Procedure UpdateLabel(node, arc, W, label, predecessor label)
    set label number l
    set label predecessor node number p
    set label predecessor label number l(p) to the label number of predecessor label
     $c(\text{node}, l) = c^{(p)} + \pi_{arc}$ 
    for all w ∈ W
        if origin of w = aerial port of node then
            add origin of node to  $L^O$ 
        end if
        for all i ∈  $L^O$ 
            if i = origin of w
                if node = destination of w then
                     $c(\text{node}, l) = c(\text{node}, l) + \alpha_w$ 
                    remove i from  $L^O$ 
                end if
            end if
        end for
    end for
end UpdateLabel

```

Figure 3-12: Label updating procedure for multilabel shortest path algorithm

3.8.4 Implementation

The Dantzig-Wolfe decomposition and column generation routines were implemented using C++ source code. The mixed integer programming formulations were solved using the callable libraries from XPress-MP (version 12). The shortest path algorithms were coded in C++ as separate functions. The graph was coded in C++ using an object oriented modeling (OOM) framework similar to the approach used by Clark [20].

3.8.4.1 Initialization

The master level scheduler must be initialized with a feasible solution. This can be accomplished using Phase I of the simplex method. However, initializing the master level scheduler with a “good” initial solution reduces the number of column generation iterations required. Specifically, a good initial solution provides a “warm start” for the decomposition and helps ensure that the dual prices passed to the subproblems are meaningful on the first iterations [8].

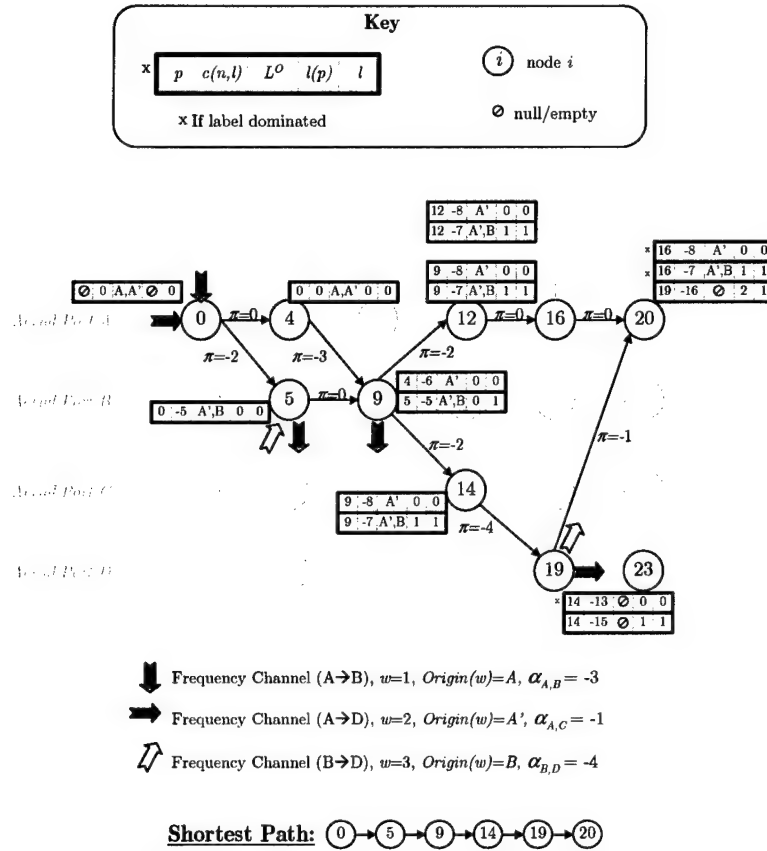


Figure 3-13: Multi-Label Shortest Path Example

There are several approaches that can be used to identify a good initial solution. Clark [20] develops a heuristic to solve a contingency munitions ship-scheduling problem. The heuristic solution is used to initialize an optimization model that improves the heuristic solution. We use a set of channel route missions from an existing real-world channel route schedule as an initial solution. The initial set of channel route missions covers all the cargo commodities and frequency requirements, which guarantees that the initial solution is a feasible solution. After scanning the initial set of channel route missions to determine the frequency channels each channel route mission services, we translate the channel route missions into aircraft route columns.

The initial set of cargo path columns is generated from the initial set of channel route missions. Each of the initial channel route missions is examined to determine the commodities for which the channel route mission is a *potential commodity cover*. A channel route mission is a potential commodity cover of a commodity if it is possible to move the commodity using the

channel route mission (assuming sufficient aircraft capacity exists). A cargo path is created for each commodity based on the channel route mission corresponding to each of the potential commodity covers. The sequence of stops for the cargo path corresponds to the sequence of stops of the channel route mission, starting at the commodity's origin and ending at the commodity's destination. An example is provided in *Example 3-5*.

Example 3-4: Consider the requirements channels, frequency channels, and the initial channel route mission below. The notation $(i,t) \rightarrow (j,t')$ indicates that an aircraft/cargo traverses the arc that originates at aerial port i in time t and terminates at aerial port j in time t' . The notation (i,j,t) indicates that a requirements channel exists between aerial port i and aerial port j and that the ALT of the cargo at aerial port i is time t .

CHANNEL ROUTE MISSION: $(A,1) \rightarrow (B,2) \rightarrow (C,3) \rightarrow (D,4) \rightarrow (A,5)$, flown by an aircraft with 50 tons of capacity.

Requirements Channels:

CHANNEL:	$(A,B,1)$	$(A,C,1)$	$(A,D,1)$	$(B,C,2)$	$(B,D,2)$
Tonnage:	20	25	30	40	20

FREQUENCY CHANNELS: (A,B) and (C,D) , both 1/week.

The following cargo paths are generated from the channel route mission:

Requirements Channel	Cargo Path Generated	Path #
$(A,B,1)$	$(A,1) \rightarrow (B,2)$	1
$(A,C,1)$	$(A,1) \rightarrow (B,2) \rightarrow (C, 3)$	2
$(A,D,1)$	$(A,1) \rightarrow (B,2) \rightarrow (C,3) \rightarrow (D, 4)$	3
$(B,C,2)$	$(B,2) \rightarrow (C, 3)$	4
$(B,D,2)$	$(B,2) \rightarrow (C,3) \rightarrow (D, 4)$	5

Table 3-4: Initial Cargo Paths Generated from an Initial Channel Route Mission

The initial set of cargo path columns and the initial aircraft route column are illustrated in Figure 3-14.

	Cargo Path 1	Cargo Path 2	Cargo Path 3	Cargo Path 4	Cargo Path 5	Aircraft Route			
	20	25	30	0	0	-50	\leq	0	(A,1)→(B,2)
	0	25	30	40	20	-50	\leq	0	(B,2)→(C,3)
	0	0	30	0	20	-50	\leq	0	(C,3)→(D,4)
	0	0	0	0	0	-50	\leq	0	(D,4)→(A,5)
	1	0	0	0	0	0	=	1	(A,B,1)
	0	1	0	0	0	0	=	1	(A,C,1)
	0	0	1	0	0	0	=	1	(A,D,1)
	0	0	0	1	0	0	=	1	(B,C,2)
	0	0	0	0	1	0	=	1	(B,D,2)
	0	0	0	0	0	1	\geq	1	(A,B)
	0	0	0	0	0	1	\geq	1	(C,D)

Figure 3-14: Initial cargo path and aircraft route columns

3.8.4.2 Column Generation Algorithm

The column generation procedure is illustrated in *Figure 3-15*. CRS-ML is solved over the initial set of columns as described in the previous section. Because the initial set of columns is restricted (i.e., it does not contain all the possible columns), CRS-ML is referred to as the *restricted master level scheduler* (CRS-RMLS). At each iteration, one subproblem is solved for each cargo commodity and each aircraft type, using dual prices from CRS-RMLS.

New columns are generated in the root node LP until a predetermined stopping criterion is met (e.g., maximum time, maximum number of CRS-RMLS iterations). To obtain an optimal integer solution, branch-and-price should be applied. In branch-and-price, the LP-relaxation at each node in a branch-and-bound tree is solved using column generation. However, we apply a heuristic technique presented by Barnhart et al. [8] in which column generation is used only at the root node in the branch-and-bound tree and the remainder of the branch-and-bound tree is solved using a standard branch-and-bound algorithm.

3.9 Results & Analysis

This section presents the results of the decomposition method applied to a small data set. The data set consists of aerial ports, aircraft, and channels in the eastern planning region (see *Figure 2-6*). The data set spans a seven-day planning horizon. Each day is split into two, twelve hour time periods. The characteristics of the data are summarized in *Table 3-5*.

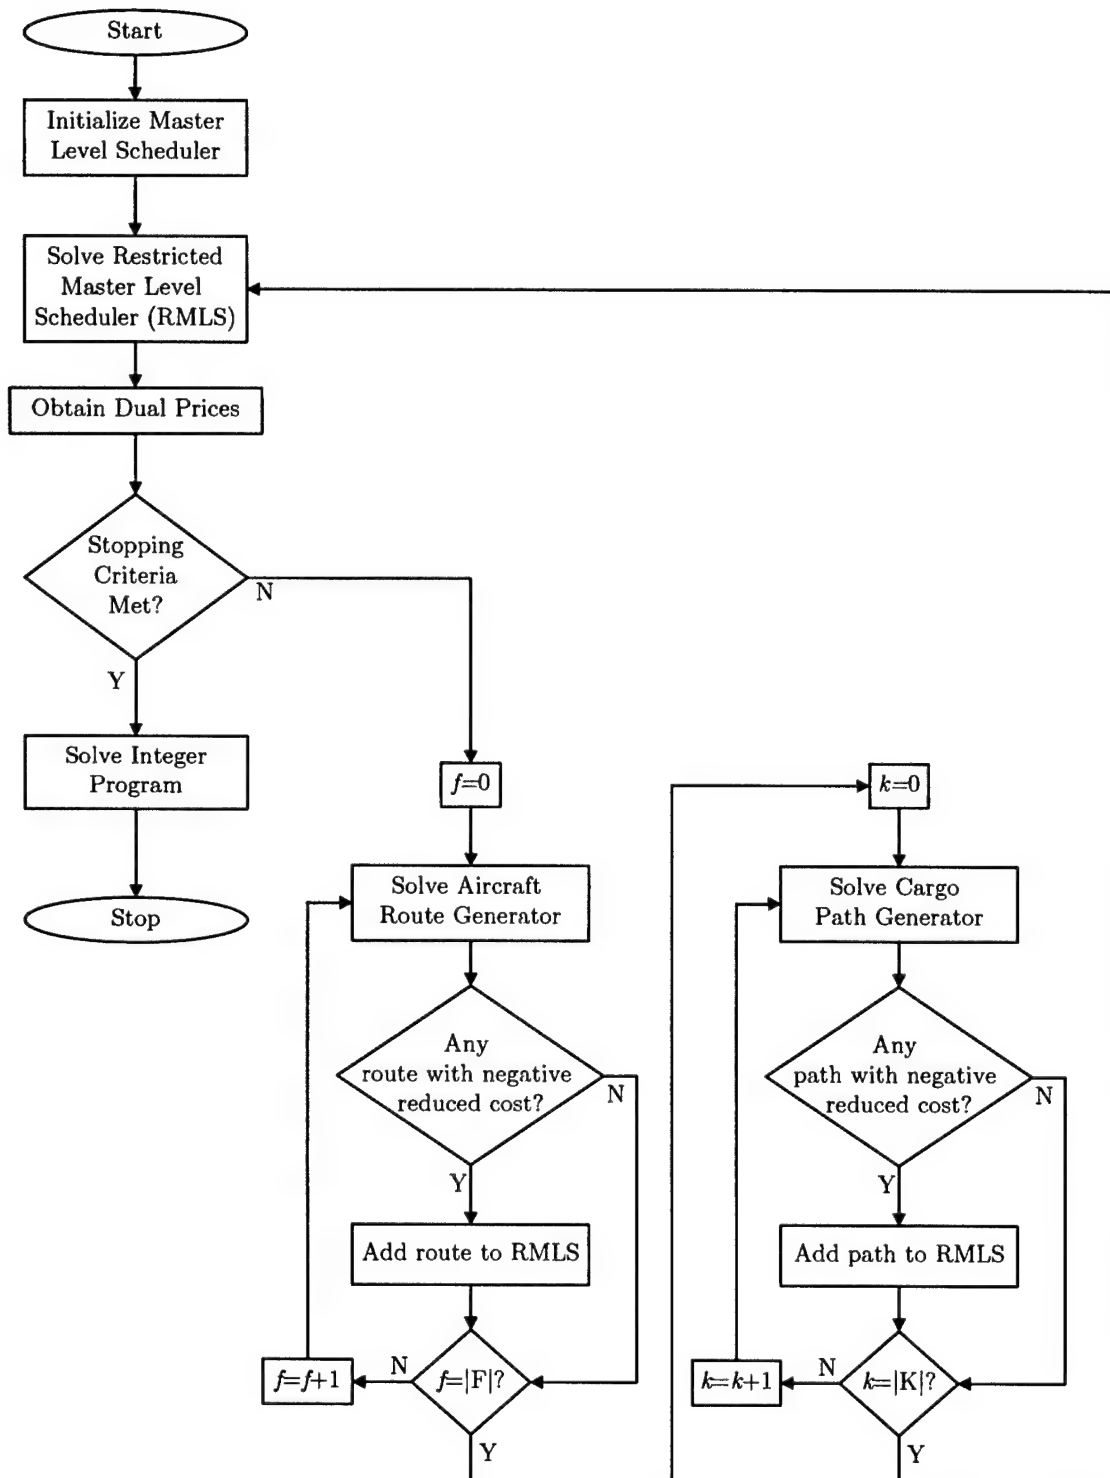


Figure 3-15: Column Generation Flow Chart

Total Aerial Ports	30
Squadron Aerial Ports	5
Planning Horizon	7 days
Periods/day	2
Commodities	254
Aircraft Types	5
Frequency Requirements	19
Fixed Buy Missions	8

Table 3-5: East Data Set Characteristics

The resulting graph contains 11,492 arcs and 420 nodes. The column generation procedure was terminated after a maximum number of iterations, and the branch-and-bound algorithm was terminated at the first integer solution. The problem was solved using a Pentium-III 550 MHZ processor and 256 megabytes of random access memory (RAM). Additionally, the pre-solve, automatic cut generation, and auto perturbation optimization solver options were turned on. The results are summarized in *Table 3-6*.

Although the path flow model (with the decomposition) is able to handle more variables than the arc flow model, the path flow model still contains forcing constraints (3.40) that cause a weak LP-relaxation. The weak LP-relaxations (i.e., large LP-IP gaps) result in the undesirable solve times in *Table 3-6*.

MAX ITERATIONS	175	200	275
<i>Rows</i>	11694	11694	11694
<i>Initial Columns</i>	7403	7403	7403
<i>Additional Columns</i>	32888	37588	49620
<i>Time to First Integer Solution (sec)</i>	5239	5386	29820
<i>Integer Solution – Number of Aircraft</i>	82	67	65
<i>Best Bound</i>	47.289	46.005	41.799
<i>LP-IP Gap</i>	42%	31%	36%

Table 3-6: Solution for Small Data set using various stopping criteria

Table 3-6 illustrates the weak LP-relaxations, especially given the presence of cuts. In addition to a weak LP-relaxation, the formulation does not accurately capture the real-world system. A more accurate representation of the real-world system requires more time periods per day. However, increasing the number of time periods per day yields a graph with significantly more arcs and nodes, which translates to more variables and constraints. Furthermore, a more accurate representation of the real-world system requires the consideration of the complex operational rules and regulations that govern the aircraft and aircrews. However, introducing these complex rules and regulations into the subproblems requires that a multilabel shortest path algorithm be used for *all* subproblems, which increases the solve time. In order to represent the real-world system more accurately, it is also necessary to enforce “rules-of-thumb” that channel route planners consider. For instance, consider the portion of the output (from the decomposition solution) presented in *Figure 3-16*. Aerial ports D, C, and N are CONUS aerial ports. The remainder of the aerial ports are OCONUS aerial ports.

The solution in *Figure 3-16* contains the routing information for the cargo commodity requiring that 0.1 tons be shipped from CONUS aerial port A, in time period 1, to OCONUS aerial port B. Notice that the cargo is shipped from the CONUS to OCONUS, back to the CONUS, and finally to its OCONUS destination. Although there is no official rule or regulation preventing this type of cargo flow, in practice channel route planners do not allow this to happen. Similar to the aircraft and aircrew rules and regulations, this type of heuristic rule would need to be enforced in the subproblems.

Overall, increasing the fidelity of the decomposition to capture the real-world system can only increase the already undesirable solution times. Therefore, we require a new formulation that can be solved quickly and that captures the fidelity required to model the real-world system.

3.10 Summary

This chapter began with a functional analysis of the channel route planning process. The functional analysis defined the appropriate inputs and outputs for the models and formulations developed in this chapter and the remainder of the thesis. Additionally, an improved operational diagram was presented that addressed the shortcomings of the current channel route planning process.

```

Column # = 46561
Commodity = A→B, in time 1 for 0.1 tons
Value in Optimal Solution = 0.1
Path: D→C→C→N→R→P→N→S→P→A→B

Aircraft traveling over Arcs in Cargo Path Column 46561:
=====
D→C : C5
C→C : Ground Arc
C→N : C5
N→R : C17
R→P : C17
P→N : C17
N→S : C17
S→P : C17
P→A : C17
A→B : C17

```

Figure 3-16: Example cargo path column from decomposition solution

Furthermore, the real-world channel route network was translated into our mathematical framework using the concept of a time-space graph. The graph was then used to develop two mathematical formulations: an arc flow network design formulation and a path flow network design formulation. The arc flow formulation did not capture all the operational rules and regulations of the real-world system, was not tractable for realistic problem instances, and had a weak LP-relaxation. The path flow formulation improved the tractability by reducing the number of constraints, at the cost of having more variables. However, the large number of variables is overcome using a decomposition approach based on Dantzig-Wolfe decomposition. By taking advantage of the underlying structure of the graph, the subproblems in the decomposition are solved quickly and efficiently to generate new columns. Although the decomposition approach improves the tractability, it still has a weak LP-relaxation. The weak LP-relaxation results in long solve times, which, in turn, prohibit the ability to increase the fidelity of the model to accurately capture the real-world system. The remainder of the thesis explores the use of a composite variable formulation, which has a stronger LP-relaxation and the ability to model transportation systems that are highly constrained by complex operational rules and regulations.

[This page intentionally left blank]

4 Composite Variable Formulation

The traditional network design formulations and the decomposition technique presented in Chapter 3 pose computational problems that cannot be overcome for realistic instances of the channel route network. In this chapter, we develop a *composite variable formulation* (CVF) that overcomes these computational challenges. The CVF has a stronger LP-relaxation, which dramatically decreases the time required to obtain an integer solution. Additionally, the CVF is able to handle complex operational rules and regulations that constrain aircraft flows and cargo flows through the channel route network.

This chapter begins with an introduction to composite variables. We discuss the motivation for using composite variables and provide a review of the literature surrounding composite variables. The idea of a CVF is to reformulate a network design formulation – such as CRM-P – solely in terms of design variables. We describe how the new design variables – *composite variables* – are constructed. Finally, we formulate and implement the channel route scheduling problem using composite variables, and compare the computational behavior of the CVF with the traditional network design formulations developed in Chapter 3.

4.1 Motivation

The formulations presented in Chapter 3 are stated in terms of cargo flow variables and aircraft route variables. This formulation strategy yields *forcing constraints* (3.28) that

typically cause fractional aircraft usage in the solution to the LP-relaxation. Furthermore, balance of flow constraints propagate fractional aircraft throughout the network, increasing the fractionality of the solution to the LP-relaxation. The result is a weak LP-relaxation, which provides a poor lower bound on the optimal integer solution. Consequently, significant effort is spent in the branch-and-bound tree to find the optimal integer solution.

The key idea of composite variables is to create a formulation that is stated solely in terms of the aircraft route variables and to implicitly capture feasible cargo flows within the aircraft route variables. The combination of aircraft route variables and cargo flow variables provide the building blocks of *composite variables*. Combining several variables to create a single variable is a common reformulation strategy. For instance, transitioning from MCNF-A to MCNF-P is a result of combining several arc flow cargo (aircraft) variables to form a single cargo (aircraft) path flow variable. The key difference between this type of reformulation and composite variables is that composite variables are based on the combination of *different* variables types – cargo flow variables and aircraft flow variables – into a single variable.

A formulation stated in terms of composite variables yields an equivalent formulation that is more computationally attractive than traditional network design formulations because of its tighter LP-relaxation [4],[22] and shorter solution times. These benefits, in theory, come at the expense of an increased number of integer variables. In practice, however, the number of variables can be limited using complex operational constraints, dominance, and branch-and-price techniques (see Cohn [22]). In general, by formulating the channel route scheduling problem using a CVF we are able to balance the complexity of the problem between the MIP formulation and the composite variable construction.

4.1.1 Composite Variable Overview

A formulation in terms of composite variables has a stronger LP-relaxation because the forcing constraints associated with traditional network design formulations can be replaced with *cover constraints*. This section illustrates this concept (adapted from Armacost [3]).

Consider the channel route network in *Figure 4-1* that consists of two aerial ports, A and B, and two aircraft types, f_1 and f_2 with capacities of 16 tons and 50 tons, respectively. Each aircraft type flies the same channel route mission, $A \rightarrow B$. There are 15 tons of cargo at aerial port A that must be delivered to aerial port B, indicated by $b_{A,B} = 15$.




Figure 4-1: Channel route network with two aerial ports and two aircraft types

There is no need to explicitly model the cargo flow variables for this channel route network because the cargo tonnage is less than the capacity of the aircraft. Rather, we use *capacity-demand constraints* to model the aircraft capacity available to *cover* the cargo commodities. The following is such a constraint.

$$16y^{f1} + 50y^{f2} \geq 15. \quad (4.1)$$

Constraint (4.1) ensures that total amount of aircraft capacity assigned to cover the cargo commodity $b_{A,B}$ is greater than its tonnage requirement. The y variables are aircraft route variables, and $y^{f1} = 1$, for example, indicates that aircraft type $f1$ flies the channel route mission $A \rightarrow B$. Relaxing the integrality restriction on the y variables and minimizing the number of aircraft will yield a fractional solution, either $y^{f1} = 15/16$ or $y^{f2} = 15/50$. Using coefficient reduction (see *Example 4-1* and Nemhauser and Wolsey [44]), the capacity-demand constraint (4.1) is equivalently written as the *cover constraint*

$$y^{f1} + y^{f2} \geq 1, \quad (4.2)$$

which removes the fractionality from the solution.

Example 4-1: The cover constraint (4.2) – transformed using coefficient reduction – removes the fractionality from the solution without discarding any feasible integer solutions, as illustrated in Figure 4-2. Region $R2$ is the feasible region for the LP-relaxation with the cover constraint (4.2) (and $y^{f1}, y^{f2} \geq 0$).




Figure 4-2: Removing fractional solutions using coefficient reduction

The region $R1 \cup R2$ comprises the feasible region for the LP-relaxation with the capacity-demand constraint (4.1) (and $y^{f1}, y^{f2} \geq 0$). The solid dots represent feasible integer solutions. Using the cover constraint (4.2) in place of the capacity-demand constraint (4.1) removes the region $R1$, which does not contain any feasible integer solutions and the LP-relaxation results in the optimal integer solution.

Because the cargo commodity's tonnage is less than the aircraft capacities, we are guaranteed that by choosing either $y^{f1} = 1$ or $y^{f2} = 1$ a feasible cargo flow for $b_{A,B}$ exists. A feasible cargo flow is not guaranteed, however, when the tonnage is greater than the capacity of at least one aircraft in the constraint. For instance, let $b_{A,B} = 30$. The capacity-demand constraint changes to

$$16y^{f1} + 50y^{f2} \geq 30. \quad (4.3)$$

Applying coefficient reduction to this constraint does not yield a cover constraint that guarantees a feasible cargo flow (i.e., $y^{f1} = 1$ does not guarantee a feasible cargo flow). Using two $f1$ aircraft will cover the tonnage, however. Thus, we create a new variable $y^{f1,f1}$, representing the combination of two $f1$ aircraft flying the channel route mission $A \rightarrow B$. The new capacity-demand constraint is stated as

$$32y^{f1,f1} + 50y^{f2} \geq 30, \quad (4.4)$$




Figure 4-3: A composite variable constructed of different aircraft types and missions

which can be transformed into a cover constraint using coefficient reduction, as done in transforming (4.1) to (4.2). The variable $y^{f1,f1}$ is a *composite variable*, and $y^{f1,f1}=1$ indicates that two $f1$ aircraft are assigned to fly the channel route mission $A \rightarrow B$.

Composite variables are not limited to combinations of aircraft of the same type flying the same channel route mission. For instance, consider the channel route network in *Figure 4-3*. The channel route network contains four aerial ports, A, B, C, and D. There are two channel route missions. Channel route mission 1 – represented by the dashed arcs in *Figure 4-3* – is one $f1$ aircraft flying $A \rightarrow B \rightarrow D$, and channel route mission 2 – represented by the solid arcs in *Figure 4-3* – is one $f2$ aircraft flying $A \rightarrow B \rightarrow C$. Independently, neither channel route mission 1 nor channel route mission 2 can completely cover all the cargo commodities incident to the mission. For instance, channel route mission 1 can only completely cover $b_{A,D}$. Similarly, channel route mission 2 can completely cover either $b_{A,B}$ or $b_{A,C}$, but not both. Combining the two missions, however, yields a composite variable $y^{f1,f2}$ that completely covers all the cargo commodities because the cargo commodity $b_{A,B}$ can be split among the two channel route missions.

We only consider composite variables that *minimally cover* a set of cargo commodities incident to the channel route missions contained in the composite variable. Minimally covering a set of cargo commodities implies that a composite variable has the following two properties.

Property 1: the composite variable completely covers a set of cargo commodities incident to the channel route missions contained in the composite variable, ensuring that feasible flows exist for each cargo commodity.

Property 2: removing one channel route mission from the composite variable does not allow the remaining channel route missions to completely cover the set of cargo commodities.

The composite variable $y^{f2,f2,f2}$ (representing three $f2$ aircraft assigned to fly channel route mission 2 in *Figure 4-3*) is not a minimal cover of $b_{A,B}$ and $b_{A,C}$ because after removing one $f2$ aircraft the remaining channel route missions can still completely cover $b_{A,B}$ and $b_{A,C}$. Property 2 prevents dominated composite variables (i.e., composite variables that will not be in the optimal solution) from being considered. For example, assuming that we are minimizing the number of aircraft, the composite variable $y^{f2,f2}$ dominates $y^{f2,f2,f2}$ because it contains fewer aircraft and completely covers the same set of cargo commodities.

The feasible minimal cover composite variables for the channel route network in *Figure 4-3* are: y^{f1} (covering $b_{A,B}$), y^{f1} (covering $b_{A,D}$), y^{f2} (covering $b_{A,B}$), y^{f2} (covering $b_{A,C}$), $y^{f1,f1,f1}$ (covering $b_{A,B}$ and $b_{A,D}$), $y^{f2,f2}$ (covering $b_{A,B}$ and $b_{A,C}$), and $y^{f1,f2}$ (covering $b_{A,B}$, $b_{A,C}$, and $b_{A,D}$). *Figure 4-4* illustrates the structure of the cover constraints that ensure all cargo commodities are covered by the composite variables where the objective is to minimize the number of aircraft.

Two key observations lead to composite variables. First, assuming that cargo is not permitted to shuffle between aircraft at intermediate locations, cargo flow variables need not be explicitly modeled in the formulation when capacity-demand constraints are used. Second, aircraft routes can be combined so that the capacity-demand constraints can be restated as cover constraints, which yield tighter LP-relaxations with superior computational behavior compared to models with capacity-demand constraints.

y^{f1}	y^{f1}	y^{f2}	y^{f2}	$y^{f1,f1,f1}$	$y^{f2,f2}$	$y^{f1,f2}$	Objective Coefficients	
1	1	1	1	3	2	2		
1	0	1	0	1	1	1	\geq	1 $b_{A,B}$
0	0	0	1	0	1	1	\geq	1 $b_{A,C}$
0	1	0	0	1	0	1	\geq	1 $b_{A,D}$

Figure 4-4: Covering constraints using composite variables

4.2 Literature Review

This section provides an overview of the composite variable literature. Although the literature focuses solely on the use of composite variables to model transportation and logistics problems, current research is exploring the use of composite variables in a variety of settings, including general network design and fixed-charge problems.

Armacost et al. [4] introduce and provide an extensive development of composite variables in the context of designing the air network for an express shipment service problem. They show the equivalence of CVFs with traditional network design problems, and they prove that the CVF achieves stronger LP-bounds on the optimal integer solution. Additionally, they show that CVFs have the ability to incorporate difficult operational constraints that render traditional network design formulations intractable. They formulate a major carrier's next day air network using a CVF, reducing the number of required aircraft by eleven percent and decreasing the total annual cost by nearly 25 percent. Armacost [3] presents composite variables in a more general setting, presenting a CVF as a Dantzig-Wolfe decomposition of a traditional network design formulation. Additionally, Armacost [3] compares composite variables to Chvátal-Gomory cuts.

Barth [13] uses the work of Armacost [3] to formulate a real-time mission planning problem in terms of composite variables. The purpose of the formulation is to generate air operations plans for a simulated wartime scenario. The formulation captures target, aircraft, armament, and routing decisions with a single composite mission variable. Composite variables implicitly capture numerous decisions and several complex constraints, allowing them to be removed from the formulation. The formulation using composite mission variables has a stronger LP-relaxation compared to conventional models used to generate air operations plans. The solution methodology employs a price-coordinated decomposition that generates composite mission variables dynamically.

Barnhart et al. [7] formulate the fleet assignment problem (see §3.2) using composite variables that represent the simultaneous assignment of fleet types to subnetworks consisting of one or more flight legs. The conventional fleet assignment model (FAM) cannot adequately capture spill costs, which occur when revenue is lost due to insufficient aircraft capacity. The itinerary based fleet assignment model (IFAM), developed by Farkas [28], accurately captures spill costs using variables that explicitly model the number of passengers spilled from each itinerary. These variables, however, introduce forcing constraints to the formulation and cause a weak LP-relaxation. Barnhart et al. [7] formulate the problem in terms of composite

variables, which accurately capture spill costs and result in an LP-relaxation that is stronger than that of IFAM.

Cohn and Barnhart [23] use composite variables to integrate the crew pairing and maintenance routing problems (see §3.2). Typically, these two problems are solved sequentially – the maintenance routing problem is solved and then aircraft crews are assigned. Crew costs are the second largest expense incurred by airlines. As a result, small improvements in the crew pairing solution can yield substantial savings. Improving the crew pairing solution, however, is limited by the set of *short connects* used in the maintenance routing solution. Unfortunately, simply solving the crew pairing problem before the maintenance routing problem does not always yield a set of short connects that is feasible for the maintenance routing problem (see Klabjan et al. [39]). On the other hand, integrating the two problems (referred to as the *basic integrated approach*) results in a large, intractable formulation (see Cordeau et al. [24]). Cohn and Barnhart [23] overcome these challenges using composite variables to implicitly capture sets of short connects that are feasible for the maintenance routing problem within crew pairing variables. Their formulation – the *Extended Crew Pairing* (ECP) model – successfully handles realistic instances and has a tighter LP-relaxation than that of the basic integrated approach.

Cohn [22] uses composite variables to formulate a service parts logistics problem. She considers a specific problem, referred to as the *high-cost, low-demand stocking problem*. Given a set of customer locations, a set of warehouses, and a set of high-cost, low-demand parts, the goal of the problem is to find the minimum-cost set of warehouses that will stock each part such that each customer location demanding a given part is covered by at least one warehouse that is assigned to stock that part. She uses a composite variable to capture the assignment of a particular *stock keeping unit* (SKU) to a *group* of warehouses. The composite variable is constructed such that the group of warehouses satisfies the covering constraints for the SKU associated with the composite variable. Defining the composite variable in this way allows the set covering constraints associated with the SKUs to be completely removed from the formulation. For large instances of the problem, the number of constraints is reduced from over a half of a million to fewer than one thousand. Reducing the number of constraints comes with a price, however. There are an exponential number of composite variables. Cohn [22] provides insights into techniques that can be used to limit the number of composite variables. Specifically, she employs the use of dominance, real-world operational rules, and a branch-and-price algorithm. Furthermore, she illustrates the importance of defining composite variables in such a way that efficient pricing problems can be constructed to identify cost-improving columns.

4.3 Formulation

Using the concepts introduced in §4.1.1, we formulate the channel route scheduling problem using composite variables. We use the following notation, in addition to the notation introduced in §3.7 (for convenience, we briefly restate below the relevant notation from §3.7), and the time-space graph introduced in §3.4.2.

SETS

PHYSICAL ASSETS IN THE CHANNEL ROUTE NETWORK (AERIAL PORTS & AIRCRAFT)

- B^S set of all aerial ports $i \in B^S$ that have an assigned aircraft squadron;
- M set of all aircraft models $m \in M$;
- F set of all aircraft types $f \in F$. Each aircraft type f is indexed by $\{m, i\}$, indicating the model type $m \in M$ and its squadron aerial port $i \in B^S$; and
- F^n set of fleet types $f \in F^n$ whose squadron base is the aerial port corresponding to node $n \in N$.

NETWORK STRUCTURE

- A set of all arcs $a \in A$ in the time-space graph;
- N set of all nodes $n \in N$ in the time-space graph;
- N^{B^S} set of squadron base nodes $n \in N^{B^S}$; and
- $N^{g(cut)}$ set of nodes $n \in N^{g(cut)}$ whose outgoing ground arc crosses the cutset time.

COMPOSITES AND AIRCRAFT MISSIONS

- FB set of commercial fixed-buy missions $l \in FB$;
- \mathbb{C}^O set of composite missions $c \in \mathbb{C}^O$ that are flown by military (i.e., organic) aircraft;
- \mathbb{C}^R set of composite missions $c \in \mathbb{C}^R$ that are flown by commercial aircraft; and
- \mathbb{C} set of composite missions $c \in \mathbb{C}$, $(\mathbb{C} = \mathbb{C}^O \cup \mathbb{C}^R)$.

COMMODITIES

- K set of all cargo commodities $k \in K$; and
- W set of all frequency channel commodities $w \in W$.

DATA

GENERAL

- \overline{H}^m upper limit on the number of required flying hours for aircraft model $m \in M$;
- \underline{H}^m lower limit on the number of required flying hours for aircraft model $m \in M$;
- ξ^w number of times frequency requirement $w \in W$ must be covered; and
- ρ^f number of available aircraft type $f \in F$.

DATA ASSOCIATED WITH COMPOSITE VARIABLES

- η_c^f the number of aircraft type $f \in F$ from composite mission $c \in \mathbb{C}$ that cross the cutset time on a flight arc;
- $\phi_c^f(n^+)$ the number of aircraft type $f \in F$ that are flying into node $n \in N$ from composite mission $c \in \mathbb{C}^O$;
- $\phi_c^f(n^-)$ the number of aircraft type $f \in F$ that are flying out of node $n \in N$ from composite mission $c \in \mathbb{C}^O$; and
- τ_c^m the number of flying hours that aircraft model $m \in M$ flies in composite mission $c \in \mathbb{C}^O$.

INDICATOR VARIABLES

- $\beta_c^w = \begin{cases} 1, & \text{if composite mission } c \in \mathbb{C} \text{ covers frequency channel } w \in W \\ 0, & \text{otherwise;} \end{cases}$
- $\delta_c^k = \begin{cases} 1, & \text{if composite mission } c \in \mathbb{C} \text{ covers cargo commodity } k \in K \\ 0, & \text{otherwise; and} \end{cases}$
- $\gamma_c^l = \begin{cases} 1, & \text{if composite mission } c \in \mathbb{C}^R \text{ corresponds to fixed-buy mission } l \in FB \\ 0, & \text{otherwise.} \end{cases}$

DECISION VARIABLES

- z_c number of composite mission $c \in \mathbb{C}$ flown;
- $g_{n^+}^f$ number of aircraft type $f \in F$ that are assigned to the ground arc incoming to $n \in N$; and

$g_{n^-}^f$ number of aircraft type $f \in F$ that are assigned to the ground arc outgoing from $n \in N$.

FORMULATION

$$\text{CRM-C} = \min \sum_{c \in \mathbb{C}^O} \sum_{f \in F} \eta_c^f z_c + \sum_{n \in N^{g(\text{end})}} \sum_{f \in F} g_{n^-}^f \quad (4.5)$$

$$s.t. \quad \sum_{c \in \mathbb{C}^O} \delta_c^k z_c + \sum_{c \in \mathbb{C}^R} \delta_c^k z_c \geq 1 \quad \forall k \in K, \quad (4.6)$$

$$\sum_{c \in \mathbb{C}^O} \beta_c^w z_c + \sum_{c \in \mathbb{C}^R} \beta_c^w z_c \geq \xi^w \quad \forall w \in W, \quad (4.7)$$

$$\sum_{c \in \mathbb{C}^O} \phi_c^f(n^+) z_c + g_{n^+}^f - \sum_{c \in \mathbb{C}^O} \phi_c^f(n^-) z_c - g_{n^-}^f = 0 \quad \forall n \in N^{B^s}, f \in F^n, \quad (4.8)$$

$$\sum_{c \in \mathbb{C}^R} \gamma_c^l z_c = 1 \quad \forall l \in FB, \quad (4.9)$$

$$\sum_{c \in \mathbb{C}^O} \eta_c^f z_c + \sum_{n \in N^{g(\text{end})}} g_{n^-}^f \leq \rho^f \quad \forall f \in F, \quad (4.10)$$

$$\sum_{c \in \mathbb{C}^O} \tau_c^m z_c \leq \overline{H}^m \quad \forall m \in M, \quad (4.11)$$

$$\sum_{c \in \mathbb{C}^O} \tau_c^m z_c \geq \underline{H}^m \quad \forall m \in M, \quad (4.12)$$

$$g_{n^+}^f \geq 0 \quad \forall n \in N, f \in F, \quad (4.13)$$

$$g_{n^-}^f \geq 0 \quad \forall n \in N, f \in F, \quad (4.14)$$

$$z_c \in \mathbb{Z}^+ \quad \forall c \in \mathbb{C}. \quad (4.15)$$

The objective function (4.5) minimizes the number of military aircraft that cross the cutset time period on either a flight arc or a ground arc. We investigate other objective functions in later chapters, demonstrating the flexibility provided by composite variables, and the impact that different objective functions have on the solution to CRM-C. *Cover constraints* (4.6) ensure that all cargo commodities are flown by either a military aircraft or a commercial aircraft. Similarly, using cover constraints (4.7), we ensure that all frequency channels are serviced by either a military aircraft or a commercial aircraft.

Constraints (4.8) ensure balance of flow for military aircraft. Balance of flow is only enforced at squadron aerial ports because operational rules require that an aircraft starts a channel route mission from its squadron base and returns to its squadron base at the conclusion of the channel route mission. We consider only the composite variables that satisfy this rule. As a result, aircraft balance of flow at non-squadron aerial ports will be satisfied without explicitly stating the associated balance of flow constraints in the formulation. Balance of flow is not required for commercial aircraft because the fixed-buy contracts only specify the routing

and timing that commercial aircraft must follow, leaving the task of positioning and depositing commercial aircraft to the commercial carriers.

There are, however, many composite variables that correspond to a single fixed-buy mission. For instance, consider a fixed-buy mission contracted for an aircraft with 90 tons of capacity to fly $A \rightarrow B \rightarrow C \rightarrow D$, and the cargo demands illustrated in *Figure 4-5*. Two composite variables are constructed for this fixed-buy mission and the set of demands because the aircraft capacity cannot completely cover all the cargo commodities. The first composite variable covers the $b_{A,B}$, $b_{A,C}$, and $b_{B,C}$ demands and the second covers the $b_{A,B}$, $b_{A,D}$, and $b_{B,C}$ demands. Each composite variable corresponds to a different way to *load* a subset of the cargo commodities on the aircraft. The subsets of cargo commodities are determined such that adding one more cargo commodity would not allow the aircraft capacity to completely cover all the cargo commodities (discussed in detail in §4.4.3). Constraints (4.9) ensure that only one of these composite variables is in the solution, capturing the fact the fixed-buy contract is for a single commercial mission that flies the route $A \rightarrow B \rightarrow C \rightarrow D$.

Aircraft count constraints (4.10) restrict the number of aircraft used to the number available. Constraints (4.11) and (4.12) ensure that each aircraft flies between the lower and upper flying hour limits. Finally, the composite variables must be integral (4.15). We only enforce that the aircraft assigned to each ground arc be non-negative (4.13)-(4.14) because constraints (4.15) and (4.8) will ensure that the ground arc variables are integral in the optimal solution.

CRM-C does not include WMOG constraints (see constraints (3.34) and (3.35) in CRM-P, §3.7.3). The purpose of CRM-C is to aid organic channel schedulers (see §2.3.3.2) in the creation of the initial cut, which does not consider WMOG limitations. Organic channel schedulers consider WMOG limitations during the execution month, only a few days prior to the actual start of each channel route mission.




Figure 4-5: Commercial composite example

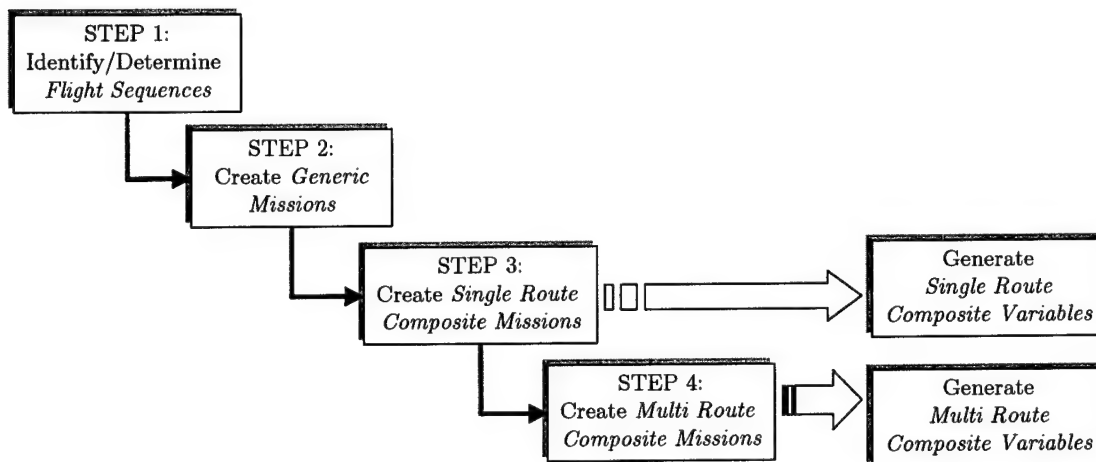


Figure 4-6: Composite Variable Generation Steps

4.4 Composite Variable Generation

CRM-C relies on the ability to identify the set of composite variables. There are many possible composite variables. Experience has illustrated, however, that real-world operational rules and regulations limit the number of feasible composite variables (see Armacost et al. [4] and Cohn and Barnhart [23]). Additionally, real-world scheduling knowledge can be built into the composite variable generation to make an initial estimate of the composite variables likely to be in an optimal solution, which further reduces the number of composite variables.

There are four steps in creating a composite variable. First, we identify a list of *flight sequences*. Second, *generic missions* are generated from the list of flight sequences. In the third step, we use the generic missions to create *single-route composite missions*. Finally, in the fourth step, we combine single-route composite missions to create *multi-route composite missions*. We create *single-route composite variables* (i.e., columns in CRM-C) after step three and *multi-route composite variables* (again, columns in CRM-C) after step four, as illustrated in Figure 4-6. In the following sections, we describe each of the four steps.

4.4.1 Step 1: Flight Sequences

The first step toward identifying the set of composite variables in CRM-C is to generate a set of *flight sequences*. A flight sequence is simply an ordered list of aerial ports that an aircraft will visit during a channel route mission. Two techniques can be used to identify flight sequences. In the first technique, the set of flight sequences is taken from an actual schedule

implemented in the real-world channel route network. Because CRM-C has a strong LP-relaxation and can be solved quickly, channel route schedulers would be able to iteratively add new flight sequences and re-solve the formulation, gaining insights about flight sequences that might improve the solution at each iteration.

The second technique enumerates the complete list of feasible flight sequences. Although the complete list of feasible flight sequences is large, heuristic scheduling rules – such as limiting the total number of aerial ports permitted in a flight sequence – reduces the number of feasible flight sequences. Political routing considerations further limits the set. For example, a country may not allow aircraft to land in certain countries prior to landing at an aerial port in their country. *Example 4-2* illustrates how flight sequences are enumerated and how various scheduling rules limit the set of flight sequences.

Example 4-2: Consider four aerial ports – A, B, C, and D – and the validated channels (A,B), (A,C), and (C,D). Aerial port A is the only squadron aerial port.

Operational rules require that an aircraft depart from its squadron aerial port at the start of a channel route mission and return to its squadron aerial port at the conclusion of the channel route mission. This rule significantly reduces the set from 45 to 15 flight sequences because a flight sequence beginning with B, C, or D is infeasible. The list of flight sequences after applying this rule is: $A \rightarrow B \rightarrow A$, $A \rightarrow C \rightarrow A$, $A \rightarrow D \rightarrow A$, $A \rightarrow B \rightarrow C \rightarrow A$, $A \rightarrow B \rightarrow D \rightarrow A$, $A \rightarrow C \rightarrow B \rightarrow A$, $A \rightarrow C \rightarrow D \rightarrow A$, $A \rightarrow D \rightarrow B \rightarrow A$, $A \rightarrow D \rightarrow C \rightarrow A$, $A \rightarrow B \rightarrow C \rightarrow D \rightarrow A$, $A \rightarrow B \rightarrow D \rightarrow C \rightarrow A$, $A \rightarrow C \rightarrow B \rightarrow D \rightarrow A$, $A \rightarrow C \rightarrow D \rightarrow B \rightarrow A$, $A \rightarrow D \rightarrow B \rightarrow C \rightarrow A$, and $A \rightarrow D \rightarrow C \rightarrow B \rightarrow A$.

Routing rules (described in §2.1.3.1) state that for each aerial port in a flight sequence, there must be a validated channel between that aerial port and another aerial port in the flight sequence. This rule reduces the set of feasible sequences to: $A \rightarrow B \rightarrow A$, $A \rightarrow C \rightarrow A$, $A \rightarrow B \rightarrow C \rightarrow A$, $A \rightarrow C \rightarrow B \rightarrow A$, $A \rightarrow C \rightarrow D \rightarrow A$, $A \rightarrow B \rightarrow C \rightarrow D \rightarrow A$, $A \rightarrow B \rightarrow D \rightarrow C \rightarrow A$, $A \rightarrow C \rightarrow D \rightarrow B \rightarrow A$, and $A \rightarrow C \rightarrow B \rightarrow D \rightarrow A$.

Consider the host country of aerial port B, and assume that the government of this country does not allow any aircraft to land in aerial port C's host country anytime prior to landing at aerial port B. Similarly, assume that the host country of aerial port C does not allow any aircraft to land in aerial port B's host country prior to landing at aerial port C. Applying this **political routing rule** limits the set of flight sequences to: $A \rightarrow B \rightarrow A$, $A \rightarrow C \rightarrow A$, and $A \rightarrow C \rightarrow D \rightarrow A$.

Although *Example 4-2* represents an extreme case, in general, the set of feasible composite variables is limited by various operational rules and regulations governing the flight sequences. When CRM-C is not solved over the entire set of composite variables (i.e., the

complete list of flight sequences is not used), the optimal solution to the original problem might not be found. However, an optimization-based approach is still attractive because the model returns an optimal solution for the *given* set of flight sequences. This approach allows the user to determine the set of flight sequences based on the desired output. For instance, because it might not be operationally feasible to change the entire channel route structure each month, this approach allows channel route planners to determine the system-wide impact of introducing a limited set of new flight sequences (e.g., redesigning a limited portion of the channel route network).

4.4.2 Step 2: Generic Missions

After the set of flight sequences is identified, we create the set of *generic missions*. A *generic mission* (GM) specifies an aircraft type and a flight sequence with generic arrival and departure times assigned to each aerial port in the flight sequence. We assume that a basic crew, which remains with the aircraft for the duration of the channel route mission, will operate each aircraft. We assign the generic timing to each aerial port in the flight sequence using aircraft and aircrew operating characteristics and regulations (e.g., flight times, crew rest times, minimum ground times). Consequently, by construction, all composite variables are feasible with respect to the complex operating rules and regulations governing aircraft and cargo flows through the channel route network. To assign generic timing to each aerial port in a flight sequence and to determine the feasibility of a generic mission, we require the following data:

FlightTime(i,j,f): the flight time between aerial port i and aerial port j for aircraft type f , which is found using the TACC Flight Time calculator.

Range(f): maximum flying range of aircraft type f .

FlightDist(i,j): the distance between aerial port i and aerial port j .

CDD(f): maximum crew duty day for a basic crew flying aircraft type f .

CR(f): minimum crew rest for a basic crew flying aircraft type f .

GrndTime(f): minimum ground time for aircraft type f .

iteration: current iteration number of *CheckFeasibility*: $iteration=0$ to $iteration=StartFreq-1$
CurrentDay: current day number (e.g., Monday = 0)
CurrentTime: time of *CurrentDay* (e.g., 1630 hrs)
HrsLastCR: hours elapsed since last crew rest
TuTAFB: total number of hours elapsed since composite route mission was initiated
ArrDay(i,f): day number that aircraft type *f* arrives at aerial port $i \in \text{flight sequence}$
ArrTime(i,f): time of *ArrDay* aircraft type *f* arrives at aerial port $i \in \text{flight sequence}$
DepDay(i,f): day number that aircraft type *f* departs from aerial port $i \in \text{flight sequence}$
DepTime(i,f): time of *DepDay* aircraft type *f* departs from aerial port $i \in \text{flight sequence}$

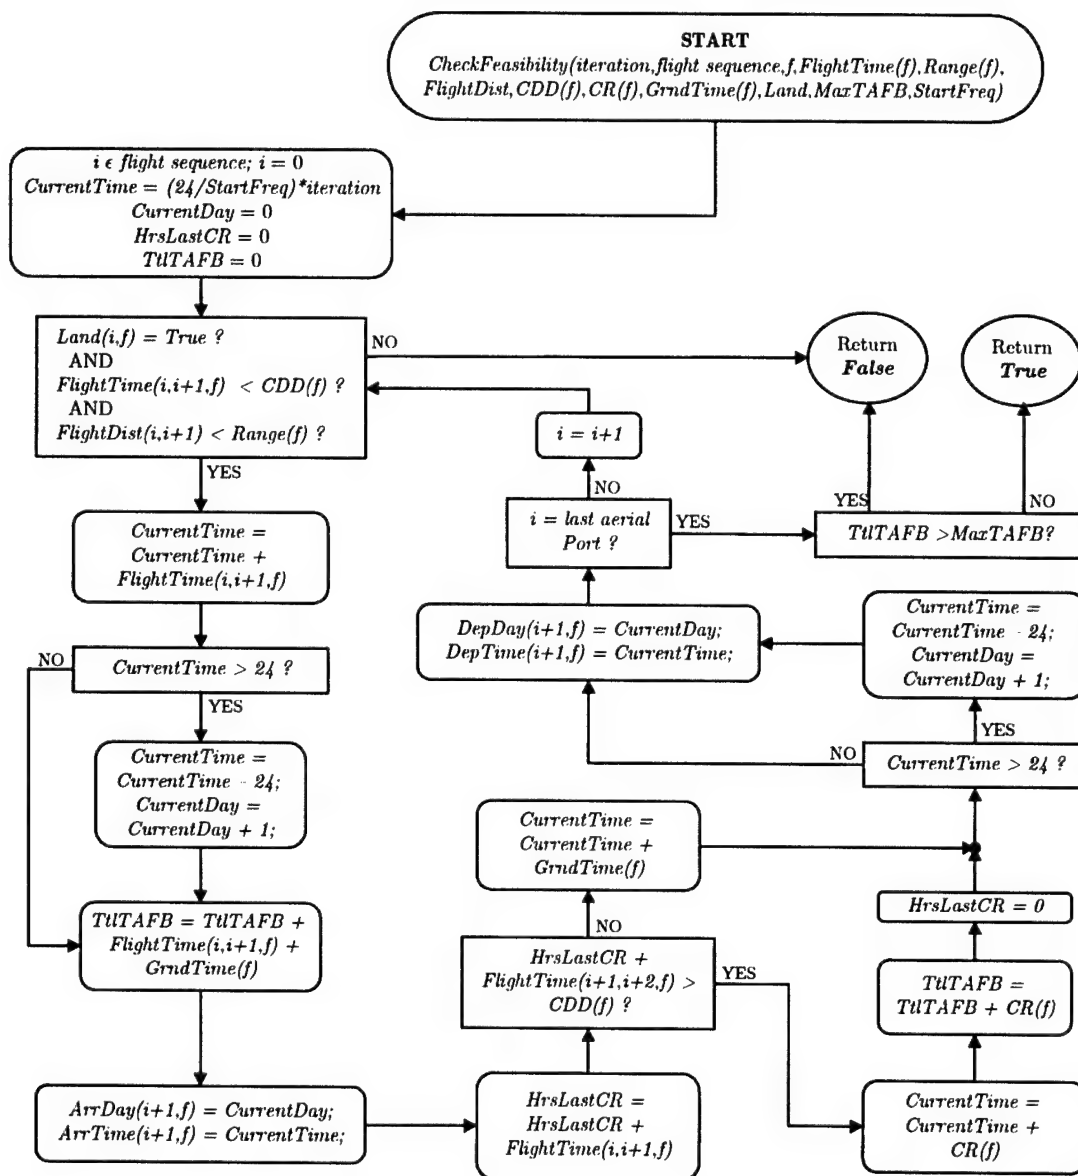


Figure 4-7: Generic Composite Mission feasibility check procedure

Land(f,i): indicates if aircraft type f is permitted to land at aerial port i .

MaxTAFB(f): parameter indicating the maximum number of hours that aircraft type f is permitted to be away from its squadron base.

StartFreq: parameter indicating the daily frequency that the generic mission is initiated at the first aerial port in the flight sequence, starting at 0000 hours. For example, given the flight sequence $A \rightarrow B \rightarrow A$, $StartFreq = 4$ indicates that the generic mission departs from aerial port A at 0000 hrs, 0600 hrs, 1200 hrs, and 1800 hrs. We define $StartFreq = 1$ to indicate that each generic mission starts at 0000 hours on each day in the planning horizon. The value of the *StartFreq* parameter corresponds to the number of time periods per day in the time-space graph.

To generate the feasible generic missions, we begin by looping through all combinations of flight sequences and aircraft types. If the aircraft type's squadron base is the first aerial port in the flight sequence, then the combination is feasible and we continue checking the feasibility of the generic mission using the procedure in *Figure 4-7*. When $StartFreq > 1$, the procedure *CheckFeasibility* (illustrated in *Figure 4-7*) is repeated $StartFreq$ times for each day in the planning horizon.

It is necessary to keep track of the time that an aircraft spends at each aerial port in the flight sequence of a generic mission. Timing information is required because when we create single-route composite missions, the day that the aircraft passes through each aerial port determines the cargo commodities and the frequency channels that the single-route composite mission can cover. Additionally, we gather data, such as total flying hours, for each generic mission. All the information associated with a generic mission is stored in data structures, making retrieval for variable construction and post processing simple. An example generic mission is illustrated in *Figure 4-8*.

4.4.3 Step 3: Single-route Composite Missions

The set of generic missions created in step two is used to generate the set of *single-route composite (SRC) missions*. An SRC mission consists of a generic mission, a specific start day in

Startday: t
 Aircraft: C5 (capacity = 50 tons); Squadron Aerial Port: A
 Flight Sequence: A→B→C→D→A

<u>AERIAL PORT</u>	<u>ARRIVALDAY</u>	<u>ARRIVALTIME</u>	<u>CR?</u>	<u>DEPARTDAY</u>	<u>DEPARTTIME</u>
A	--	--	--	t	0000
B	t	1300	yes	t+1	0500
C	t+1	1230	no	t+1	1545
D	t+1	2045	yes	t+2	1330
A	t+2	2300	--	--	--

Flying Hours: 35.75 hours
 Total TAFB: 71.00 hours

Figure 4-8: Example of a generic mission

the planning horizon, a set of *potential* cargo commodities that can be covered, a set of *actual* cargo commodities covered, and a set of frequency channels covered.

We begin generating SRC missions by looping through all combinations of generic missions and days in the planning horizon. The generic timing assigned to the flight sequence in the generic mission is updated to reflect the starting day of the SRC mission. We maintain a list of operating hours, BASH and quiet hours, as well as base closures for each aerial port on each day in the planning horizon. The list is used to test the feasibility of each generic mission-day combination, ensuring that each aerial port in the flight sequence – updated to reflect the starting day – is open for operation. If each aerial port is open, we scan the SRC mission's flight sequence to determine the cargo commodities and frequency channels that can potentially be covered by the SRC mission. A cargo commodity can potentially be covered by a SRC mission if both the origin and destination of the cargo commodity are contained in the flight sequence, and if the aircraft departs the cargo commodity's origin within PHT time periods of the cargo commodity's available to load time. Similarly, a frequency channel is covered by a SRC mission if both the origin and destination of the frequency channel are contained in the flight sequence, and if the aircraft arrives at the frequency channel's destination within the frequency of service requirement.

Each cargo commodity that can potentially be covered by the SRC mission is added to a list of *potential cargo commodities*, which contains all the cargo commodities that could be moved by an aircraft over the flight sequence assuming unlimited aircraft capacity. It should be noted that, by construction, each cargo commodity in the potential cargo commodity list of a SRC mission satisfies the PHT service level requirement. Each frequency channel that can be serviced by the SRC mission is added to the list of frequency channels covered by the SRC mission.

Next, we determine which of the potential cargo commodities can be completely covered by the SRC mission, given the capacity of the generic mission aircraft. There are four possible outcomes, which we describe below.

Case 1: *Sufficient aircraft capacity exists to completely cover all potential cargo commodities.*

In this case, each cargo commodity is removed from the potential list of cargo commodities and added to the list of *actual* cargo commodities covered. We create a *SRC variable* from the SRC mission.

Each SRC variable is represented as a column in CRM-C. We use the list of actual cargo commodities covered, the list of frequency channels covered, and the departure and arrival times at the squadron aerial port of the assigned aircraft to determine the column's constraint coefficients (e.g., $\beta_c^w, \delta_c^k, \phi_c^f$).

Case 2: *Sufficient aircraft capacity exists to cover all of the cargo commodities in the list of potential cargo commodities, after each cargo commodity with a tonnage requirement greater than the aircraft capacity is removed from the list of potential cargo commodities.*

If the aircraft capacity cannot completely cover all the cargo commodities in the potential cargo commodity list of an SRC mission, the list of potential cargo commodities is scanned. Each cargo commodity with a tonnage requirement greater than the aircraft capacity is removed from the potential list and added to a list of *non-covered* cargo commodities. After removing a cargo commodity from the list of potential cargo commodities, we refer to the list as the *updated potential cargo commodity list*. Information is kept indicating that the SRC mission has residual aircraft capacity, noted \hat{u}_*^k (where $*$ is the SRC mission number), to possibly transport the cargo commodity added to the non-covered list. If all of the cargo commodities in the updated potential cargo commodity list can be completely covered, a single-route composite variable – covering the cargo commodities in the updated potential cargo commodity list – is created, as described in Case 1 above. We provide an example below.

Example 4-3: *Consider the network in Figure 4-9 and an aircraft with 50 tons of capacity. SRC mission #1 is depicted by the dashed arcs: $A \rightarrow B \rightarrow C \rightarrow D \rightarrow A$. The list of potential cargo commodities for SRC mission #1 is $\{b_{A,B}, b_{A,C}, b_{A,D}, b_{B,C}\}$. The tonnage of cargo commodity $b_{A,B} = 80$ is greater than the 50 ton aircraft capacity. We remove $b_{A,B}$ from the list of potential cargo commodities and add it to the non-*




Figure 4-9: Network for Single-route Composite Mission - Case 2

covered list, noting the residual aircraft capacity available $\hat{u}_1^{b_{A,B}} = 15$. All remaining cargo commodities in the updated potential cargo commodity list can be completely covered, so each cargo commodity is removed and added to the list of actual cargo commodities covered, which is $\{b_{A,C}, b_{A,D}, b_{B,C}\}$.

SRC Mission #2 is depicted by the solid arcs: $A \rightarrow B \rightarrow E \rightarrow A$. The list of potential cargo commodities for SRC mission #2 is $\{b_{A,B}, b_{A,E}, b_{B,E}\}$. The cargo commodity $b_{B,E} = 80$ is greater than the 50 ton aircraft capacity, so it is removed from the list of potential cargo commodities and added to the non-covered list, noting the residual aircraft capacity available $\hat{u}_2^{b_{B,E}} = 20$. Additionally, $b_{A,B}$ cannot be completely covered, so it is removed from the potential cargo commodity list of SRC mission #2. Because cargo commodity $b_{A,B}$ is already in the non-covered list, we simply add SRC mission #2 with residual $\hat{u}_2^{b_{A,B}} = 20$ to the list of other missions (i.e., SRC mission #1) that have residual capacity available for cargo commodity $b_{A,B}$. We keep this information so that eventually we can piece together residual aircraft capacities from various SRC missions such that the combined residual capacities completely cover the non-covered cargo commodity $b_{A,B}$. The list of actual cargo commodities covered by SRC mission #2 is $\{b_{A,E}\}$. Table 4-1 summarizes the information associated with each SRC mission.

SRC Mission #	Potential Cargo Commodity List	Actual Cargo Commodity List	Non-covered Commodities (Commodity, \hat{u}_i^t)
1	$\{b_{A,B}, b_{A,C}, b_{A,D}, b_{B,C}\}$	$\{b_{A,C}, b_{A,D}, b_{B,C}\}$	$(b_{A,B}, \hat{u}_1^{b_{A,B}} = 15)$
2	$\{b_{A,B}, b_{A,E}, b_{B,E}\}$	$\{b_{A,E}\}$	$(b_{A,B}, \hat{u}_2^{b_{A,B}} = 20)$ $(b_{B,E}, \hat{u}_2^{b_{B,E}} = 20)$

Table 4-1: Single-Route Composite Missions

Case 3: Sufficient aircraft capacity exists to cover some of the cargo commodities in the list of potential cargo commodities, after each cargo commodity with a tonnage requirement greater than the aircraft capacity is removed from the list of potential cargo commodities.

Case 3 is similar to case two, except that after removing all the cargo commodities with a tonnage requirement greater than the aircraft capacity from the list of potential cargo commodities, the aircraft capacity cannot cover *all* the cargo commodities in the updated potential cargo commodity list. As in Case 2, we keep track of the residual aircraft capacity available for the cargo commodities removed from the list of potential cargo commodities.

We use the updated potential cargo commodity list to identify the set of *maximal cover cargo loadings*. A *cargo loading* is simply a subset of the cargo commodities in the updated potential cargo commodity list that can be completely covered by the aircraft capacity. A cargo loading is a *maximal cover* if adding one more cargo commodity will exceed the aircraft capacity on at least one flight leg in the flight sequence. We only consider maximal cover cargo loadings to prevent dominated SRC missions from being considered. There is, however, a potential drawback – which we discuss in §4.4.4 – of considering only maximal cover cargo loadings. A SRC variable is generated – as described in Case 1 above – for each maximal cover cargo loading. *Example 4-4* illustrates Case 3.

Example 4-4: Consider the network in Figure 4-10 and an aircraft with 50 tons of capacity. Consider a SRC mission with the flight sequence depicted by the solid arcs: $A \rightarrow B \rightarrow C \rightarrow D \rightarrow A$. The list of potential cargo commodities is $\{b_{A,B}, b_{A,C}, b_{A,D}, b_{B,C}, b_{B,D}, b_{C,D}\}$. After $b_{A,B}$ is removed (because its tonnage requirement is greater than the aircraft capacity), the SRC mission cannot completely cover all the cargo commodities in the updated potential cargo commodity list. Table 4-2 contains the SRC missions generated for the maximal cover cargo loadings.




Figure 4-10: Network for Single-route Composite Mission with multiple cargo loadings

SRC Mission #	List of Actual Cargo Commods (i.e., maximal cover cargo loading)	Additions to Non-covered list (Commodity, \hat{u}_i^t)
3	$\{b_{A,C}, b_{B,C}, b_{C,D}\}$	$(b_{A,B}, \hat{u}_3^{b_{A,B}} = 15)$
4	$\{b_{A,D}, b_{B,C}, b_{B,D}\}$	$(b_{A,B}, \hat{u}_4^{b_{A,B}} = 30)$
5	$\{b_{B,C}, b_{B,D}, b_{C,D}\}$	$(b_{A,B}, \hat{u}_5^{b_{A,B}} = 50)$

Table 4-2: Single-route Composite Missions with multiple cargo loadings

Case 4: Sufficient aircraft capacity does not exist to cover any of the cargo commodities in the list of potential cargo commodities.

In this case, each of the cargo commodities in the list of potential cargo commodities has a tonnage requirement greater than the aircraft capacity. Each cargo commodity is removed from the list of potential cargo commodities and added to the list of non-covered commodities, noting that the SRC mission has residual aircraft capacity to cover the cargo commodity added to the non-covered list. The residual capacity available is the entire aircraft capacity and will be used – in combinations with residual aircraft capacities from other SRC missions – to cover the non-covered cargo commodities. No SRC variable is generated for the formulation. *Example 4-5* illustrates case four.

Example 4-5: Consider the network in Figure 4-11, an aircraft with 50 tons of capacity, and SRC mission #6 depicted by the dashed arcs: $A \rightarrow B \rightarrow A$. The list of potential cargo commodities is $\{b_{A,B}\}$, which cannot be completely covered. The cargo commodity $b_{A,B}$ is added to the list of non-covered cargo commodities, noting the residual aircraft capacity available $\hat{u}_6^{b_{A,B}} = 50$.




Figure 4-11: Example Network for Single-route Composite Mission generation

4.4.4 Step 4: Multi-route Composite Missions

For each cargo commodity in the non-covered cargo commodity list we create minimal cover (see §4.1.1) *multi-route composite (MRC) mission*. A MRC mission is composed of two or more SRC missions, a list of actual cargo commodities covered, and a list of covered frequency channels. A MRC variable is generated for each MRC mission.

Similar to SRC variables, a MRC variable is represented as a column in CRM-C. We use the list of actual cargo commodities covered, the list of frequency channels covered, and the departure and arrival times at the squadron aerial port of the assigned aircraft from each SRC mission in the MRC mission to determine the correct constraint coefficients of the column (e.g., $\beta_c^w, \delta_c^k, \phi_c^f$). *Example 4-6* illustrates how we create MRC missions.

Example 4-6: Recall SRC missions 3, 4, 5, and 6 from *Example 4-4* and *Example 4-5*, whose flight sequences are shown in *Figure 4-12*. The non-covered cargo commodity list from *Examples 4-4* and *4-5* contains the cargo commodity $b_{A,B}$. SRC missions 3, 4, 5, and 6 have residual aircraft capacity available for cargo commodity $b_{A,B}$. We piece together enough residual capacity from SRC missions 3, 4, 5, and 6 to minimally cover the cargo commodity in the list of non-covered cargo commodities (i.e., cargo commodity $b_{A,B}$). Three feasible MRC missions are contained in *Table 4-3*.

MRC Mission #	SRC Missions $\{SRC \text{ Mission \#, Residual}\}$	Actual Cargo Commodities Covered
1	$\{4, \hat{u}_4^{b_{A,B}} = 30\}, \{6, \hat{u}_6^{b_{A,B}} = 50\}$	$\{b_{A,B}, b_{A,D}, b_{B,C}, b_{B,D}\}$
2	$\{5, \hat{u}_5^{b_{A,B}} = 50\}, \{6, \hat{u}_6^{b_{A,B}} = 50\}$	$\{b_{A,B}, b_{B,C}, b_{B,D}, b_{C,D}\}$
3	$\{4, \hat{u}_4^{b_{A,B}} = 30\}, \{5, \hat{u}_5^{b_{A,B}} = 50\}$	$\{b_{A,B}, b_{A,D}, b_{B,C}, b_{B,D}, b_{C,D}\}$

Table 4-3: Multi-Route Composite Missions




Figure 4-12: Example Network for Multi-route Composite Mission

The list of actual cargo commodities covered by each MRC mission contains the cargo commodities covered by each SRC mission and the cargo commodity the MRC mission is constructed to cover. For instance, the list of actual cargo commodities covered for SRC mission #4 is $\{b_{A,D}, b_{B,C}, b_{B,D}\}$, and the list of actual cargo commodities covered for SRC mission #6 is $\{\emptyset\}$. These two SRC missions have residual aircraft capacity to cover the non-covered cargo commodity $b_{A,B}$, and are combined to create MRC mission #1. The list of actual cargo commodities covered by MRC mission #1 is $\{b_{A,B}, b_{B,C}, b_{B,D}, b_{C,D}\} = \{b_{A,D}, b_{B,C}, b_{B,D}\} \cup \{\emptyset\} \cup \{b_{A,B}\}$.

It is possible, as in MRC mission #3 in Table 4-3, that the combined SRC missions cover one or more of the same cargo commodities. Overlapping cargo commodities are removed from one of the SRC missions and the increased aircraft capacity is used as residual capacity for the cargo commodity that the MRC mission is being constructed to cover.

If two SRC missions service the same frequency channel, however, the MRC mission is given credit in CRM-C for covering the frequency channel multiple times. For example, assume that both of the SRC missions in MRC mission #3 service a frequency channel (A,D). MRC mission #3, in turn, services the frequency channel twice.

In this chapter, we assume that fixed-buy SRC missions cannot be combined to create multi-route composite missions. The justification is that it is operationally difficult to split cargo commodities either between a commercial aircraft and a military aircraft or between several different commercial aircraft types (e.g., some commercial aircraft require *contoured cargo pallets*, which means that cargo pallets must be specially built so that the cargo on the pallet fits to the contour of the aircraft's cargo bay).

As described in §4.4.3, we only consider maximal cover cargo loadings. This has an impact on our ability to successfully construct MRC missions that completely cover the cargo commodities in the non-covered list. Consider the MRC missions in Table 4-3. We were able to construct these MRC missions because the residual aircraft capacity of each SRC mission is relatively large. Considering only maximal cover cargo loadings might prevent SRC missions with enough residual aircraft capacity from being generated, however. For instance, if SRC mission #3 ($\hat{u}_3^{b_{A,B}} = 15$) and SRC mission #4 ($\hat{u}_4^{b_{A,B}} = 30$) from Table 4-2 were the only candidate SRC missions for cargo commodity $b_{A,B} = 80$, there would not be enough residual aircraft capacity available to cover $b_{A,B}$. For the data sets we've considered, this does not occur. For instances in which the data does not prevent this from happening, the requirement of maximal cover cargo loadings is relaxed. We generate SRC missions that have dominated cargo loadings,

but we do not generate SRC mission variables from these missions. Rather, these missions are used only to construct MRC missions.

4.5 Composite Variable Generation Considerations

Aircraft capacities, the number of flight sequences, the maximum number of SRC missions that can be combined to create an MRC mission, and the number of time periods per day all influence composite variable generation.

AIRCRAFT CAPACITIES

Aircraft capacities influence the cargo commodities that can be covered by a SRC mission. The maximum aircraft takeoff weight for each aircraft type presented in §2.2.3 is the maximum allowed combined weight of fuel and cargo. An aircraft must carry more fuel the farther it flies, reducing the amount of cargo that can be carried. The relationship between distance and available cargo capacity is described by each aircraft type's *range-payload curve*. The range-payload curve determines the amount of cargo capacity that is available, given the distance the aircraft must fly (see *Figure 4-13*).

Although we assume constant aircraft capacity, the structure of CRM-C is not influenced by the range-payload curve (see Armacost [3]). The range-payload curve would be used as an input to the composite variable generation procedures. When an aircraft is assigned to a flight sequence during composite variable generation, the available aircraft capacity would be determined from the range-payload curve and the distance flown (determined from the flight sequence which is known a priori). The remainder of the composite variable generation would continue as described above.




Figure 4-13: Notional range-payload curve

MULTI-ROUTE COMPOSITE MISSION CONSTRUCTION

Operationally, splitting cargo commodities between many aircraft might not be desirable. One reason this might not be desirable, for instance, is that the aerial ports attempt to build entire pallets that contain cargo for a common destination. Thus, assigning significantly less than a pallet load of cargo to flow on an aircraft either underutilizes the aircraft or significantly increases the required ground time for an aircraft. For instance, consider an aircraft traveling $A \rightarrow B \rightarrow C$ that is carrying a cargo pallet consisting of cargo destined for both aerial port B and aerial port C. At aerial port B, the entire pallet must be offloaded from the aircraft and partially unpacked. Then, the pallet must be reloaded onto the aircraft so that the remaining cargo on the pallet can continue to aerial port C.

We parameterize the amount of residual capacity that an SRC mission must have available for it to be considered as a candidate for inclusion in an MRC mission. For instance, consider a 150-ton cargo commodity. During MRC mission generation, we only consider the SRC missions that have a residual capacity of at least a predetermined fraction of the 150 tons. The size of the fraction is set to prevent cargo commodities from being split into small pieces. The size of the fraction can be changed and differ for each cargo commodity. The size of the fraction influences the solution because a larger fraction generally results in MRC missions that are combinations of large capacity aircraft. We test the model's sensitivity to this parameter in Chapter 5.

Similarly, SRC missions that are candidates for inclusion in an MRC mission can visit the origin aerial port of the same cargo commodity on different days, requiring the cargo commodity to be split between multiple aircraft and split over several days. *Example 4-7* illustrates this idea.




Figure 4-14: Example Network for Multi-route Composite Mission generation

Example 4-7: Consider the graph in Figure 4-14, which illustrates two SRC missions. SRC mission #1 is depicted by the dashed arcs and SRC mission #2 is depicted by the solid arcs. Assume that the two SRC missions are being combined to create a MRC mission to cover the cargo commodity $b_{A,B}$. Combining these two missions requires that the cargo commodity be split over two time periods. A fraction of $b_{A,B}$ is flown on SRC mission #1, departing in time period one, and the remainder of $b_{A,B}$ is flown on SRC mission #2, which departs in time period three.

Initially, we do not allow SRC missions that require cargo commodities to be split over more than one time period to be combined to create an MRC mission. We do, however, treat this as a parameter and investigate different time intervals in Chapter 5. Other than affecting the number of possible MRC missions, allowing different time intervals does not change the composite variable generation procedures described in this chapter.

INFLUENCE OF THE *STARTFREQ* PARAMETER ON TIME

The *StartFreq* parameter, defined in §4.4.2, plays a key role in defining the size of CRM-C, influencing both the number of constraints and the number of composite variables. Larger values of *StartFreq* increase the number of nodes in the graph and, consequently, the number of balance of flow constraints (4.8). Larger values of *StartFreq* also increase the number of SRC variables because there are more starting times for each generic mission. More SRC variables, in turn, increase the number of MRC variables. Initially, we set *StartFreq* equal to one. We discuss various values of *StartFreq* in Chapter 5.

When *StartFreq* equals one, we assume that each SRC mission starts at 0000 hours on each day in the planning horizon. This poses problems for the aircraft balance of flow at squadron aerial ports, which we illustrate in *Example 4-8*.




Figure 4-15: Influence of *StartFreq* on aircraft balance of flow

Example 4-8: Consider the graph in Figure 4-15a, which depicts two SRC missions. SRC mission #1 flies $A \rightarrow B \rightarrow A$, starting at 0000 hours on day 1 (depicted by the dashed arcs), and SRC mission #2 flies $A \rightarrow C \rightarrow B \rightarrow A$, starting at 0000 hrs on day 3 (depicted by the solid arcs). Suppose that SRC mission #1 returns to aerial port A at 1230 hours on day 3. As modeled in Figure 4-15a, the same aircraft assigned to fly SRC mission #1 can also be assigned to fly SRC mission #2. This is not feasible, however, because SRC mission #1 does not arrive at aerial port A on day 3 until after SRC mission #2 is scheduled to depart. To prevent the same aircraft from being assigned to both missions, we do not allow the aircraft assigned to SRC mission #1 to be assigned to another SRC mission until the next time period, which is 0000 hours on day 4 (illustrated in Figure 4-15b). We call this the *ready time*. The ready time only influences the balance of flow constraints (4.8) in CRM-C, not the composite route generation procedures described in this chapter.

INFLUENCE OF THE *STARTFREQ* PARAMETER ON SLACK

A viable schedule does not follow immediately from the output of CRM-C. There is *slack* in the solution, allowing channel route schedulers to appropriately sequence aircraft into each aerial port so that WMOG limitations are observed. Specifically, the slack represents the amount of time that the departure time of each SRC mission can be modified without either changing the cargo commodities covered, the frequency channels covered (which are ultimately based on the assumed departure time dictated by *StartFreq*), or the ability of the aircraft to arrive at each aerial port when it is open for operations. The *StartFreq* parameter influences the amount of slack in the solution of CRM-C. We illustrate the concept of slack in *Example 4-9*.

Example 4-9: Consider the following SRC composite mission:

Startday: 4
Aircraft: C5; capacity = 50 tons
Squadron Aerial Port: A
Flight Sequence: $A \rightarrow B \rightarrow C \rightarrow D \rightarrow A$

AERIAL PORT	ARRIVALDAY	ARRIVALTIME	CR?	DEPARTDAY	DEPARTTIME
A	--	--	--	4	0000
B	4	1300	yes	5	0500
C	5	1230	no	5	1545
D	5	2045	yes	6	1330
A	6	2300	--	--	--

Flying Hours: 35.75 hours
Total TAFB: 71.00 hours

Assume that the list of actual cargo commodities covered by the SRC mission (with the available to load times (ALT) indicated by day) is $\{(b_{A,B}, ALT=4), (b_{A,C}, ALT=4), (b_{B,C}, ALT=2), (b_{C,D}, ALT=4)\}$. It is important to note that the actual list of cargo commodities covered is determined from the potential list of

cargo commodities, which is based on the departure time dictated by the *StartFreq* parameter (e.g., 0000 hrs). Assume that the operating hours (normalized to zulu time) for the aerial ports are:

AERIAL PORT	OPEN	CLOSE
A	24 hrs/day	
B	0600	1630
C	1100	2330
D	0930	2330

The slack for the above SRC mission is two hours, indicating that the *DepartTime* at aerial port A on day 4 can be anytime between 0000 hours and 0200 hours. Anytime later than 0200 hours causes the aircraft to arrive at a downstream aerial port during closed hours, or hinders the ability of the SRC mission to cover its cargo commodities and/or frequency channels. For instance, a departure time of 0300 hours from aerial port A puts the aircraft at aerial port D on day 6 at 0030 hours, which is not feasible because aerial port D is closed.

Assume that the maximum port hold time permitted for each cargo commodity is two days. If the SRC mission is permitted to start at aerial port A at 1400 hours on day 4, the aircraft will arrive at each downstream aerial port during operating hours. The aircraft will arrive, however, at aerial port B on day 5, which means that the $(b_{B,C}, \text{ALT}=2)$ cargo commodity must remain in port for 3 days, exceeding the maximum port hold time of two days.

There is an important connection between *StartFreq*, slack, and the ability of channel route planners to transition the solution of CRM-C into a channel route schedule that is operationally feasible. Given a value for *StartFreq*, the slack is defined by the aerial port operating hours and the cargo commodities and frequency channels covered. The amount of slack limits the amount of time that a SRC mission's start time can be shifted, which can be overly restrictive. Consider *Example 4-9*. A slack of two hours at a start time of 0000 hours might not accurately capture when channel route missions typically depart. Redefining the start time associated with *StartFreq*=1 (e.g., changing *StartFreq*=1 to indicate that each generic mission starts at 1200 hours on each day in the planning horizon) can alleviate this problem. Additionally, increasing the value of *StartFreq* captures various start times (i.e., more fidelity), independent of how *StartFreq*=1 is defined. The drawback to larger values of *StartFreq*, however, is an increased number of variables. There is a balance between the appropriate level of fidelity and the model size determined by *StartFreq*, which we discuss in Chapter 5.

4.6 Results & Analysis

This section presents the results of CRM-C applied to the small set of aerial ports and aircraft used in Chapter 3. As a result of CRM-C's superior computational behavior, however, we are able to increase the planning horizon to 30 days. The characteristics of the data set are summarized in *Table 4-4*.

The purpose of this section is to illustrate the computational behavior of CRM-C, not to analyze the solution, which we do in Chapter 5. We implement a special case of CRM-C and compare the computational behavior to the models presented in Chapter 3. As in Chapter 3, we relax the flying hour constraints (4.11) and (4.12).

Total Aerial Ports	30
Squadron Aerial Ports	5
Planning Horizon	30 days
Cargo Commodities	1016
Aircraft Types	5
Frequency Channels	67
Fixed-buy Missions	32

Table 4-4: Data set characteristics

We use the same set of flight sequences – taken from an existing real-world channel route schedule – used to initialize CRM-P in Chapter 3. We set the maximum port hold time to two days. We set the maximum time away from base high enough to ensure that all generic missions are feasible. We set *StartFreq*=1, and we require that a SRC mission have at least twenty percent of a cargo commodity's tonnage of residual capacity available to be considered as a candidate for an MRC mission. For instance, for a 120-ton cargo commodity we require that an SRC mission have at least 24 tons of residual capacity available. Furthermore, we require that each aircraft in an MRC mission depart the origin aerial port of the cargo commodity that the MRC mission is constructed to cover within one day of the other aircraft in the MRC mission.

The formulation was solved using a Pentium-III 550 MHZ processor and 256 megabytes of RAM. The composite variable generation routines were implemented in C++ and the MIP formulations were solved using the callable libraries from XPress-MP (version 12). The computational results, which we refer to as the *baseline* CRM-C results, are summarized in *Table 4-5*.

COMPOSITES	Flight Sequences	139
	Fixed-buy	3314
	Single-route	7954
	Multi-route	5534
PROBLEM SIZE	Columns	16982
	Rows	1301
	Non-zeroes	197596
OBJECTIVE VALUE <i>Number of Aircraft</i>	LP Relaxation	21.96879
	First Integer	23
	LP-IP Gap	0.046
RUN TIME (SEC.)	LP Relaxation	317
	First Integer	1596

Table 4-5: Computational Results of baseline CRM-C solution

The results in *Table 4-5* illustrate the strength of the LP-relaxation of CRM-C, especially considering the increased size of the data set resulting from the increased planning horizon. Furthermore, by carefully constructing the composite variables, we ensure that all the composite variables in the solution satisfy the complex operational rules and regulations that the formulations developed in Chapter 3 could not easily capture.

4.7 Summary

This chapter began with an introduction to composite variables, which we use for two reasons to model the channel route network. First, the LP-relaxation of traditional network design formulations are weak, making realistic instances of the channel route network difficult to solve. Second, complex rules and regulations that constrain aircraft and cargo flows through the channel route network are difficult to enforce in traditional network design formulations. Composite variables overcome both of these shortcomings because the forcing constraints that cause the poor LP-relaxations in traditional network design formulations can be replaced by cover constraints, and the composite variable generation procedures enforce complex operational rules.

We formulate the channel route scheduling problem using a composite variable formulation, and implement the formulation using a small data set. The results illustrate the improved computational behavior of CRM-C compared to the formulations developed in Chapter 3.

[This Page Intentionally Left Blank]

5 Results and Analysis

In Chapter 4, we formulated the channel route scheduling problem using composite variables, and we implemented the formulation using a sample data set. We used the results to illustrate the composite variable formulation's improved computational behavior compared to traditional network design formulations. In this chapter, we extend our analysis, examining both the solution quality and tractability of CRM-C using different objective functions, and different values of the parameters described in the Chapter 4.

The current channel route planning process is largely manual, and time consuming. This prevents channel route planners from investigating different route structures. Additionally, there is no formal method available in the current channel route planning process to quantify the system-wide impact of changes to the channel route network, such as changes in cargo demand levels, aircraft availability, and route structure. One purpose of this chapter is to illustrate the ability of CRM-C to overcome these shortcomings.

We begin the chapter with a description of the metrics we use to evaluate the quality of the solution to CRM-C. Next, we describe the objective functions, and we emphasize how both the metrics and objective functions relate to either the current channel route planning process or one of the three operational objectives introduced in Chapter 2. We decompose our analysis in this chapter into two parts. First, we investigate the tractability and size of CRM-C using the different objective functions and different parameter values. Second, we use the metrics introduced at the beginning of the chapter to investigate the characteristics of the solution to

CRM-C when it is solved using different objective functions and different parameter values. We connect the results of the two types of analysis by illustrating that there are large instances of CRM-C – which result from changing certain parameter values – that do not need to be considered because the parameter values do not lead to improved solution quality.

The range of analysis in this chapter highlights a key attribute of composite variable formulations – the structure of the formulation does not change as we evaluate different scenarios. The changes occur, rather, in the composite variable generation procedures that we developed in Chapter 4. This is attractive because it allows the model to be easily manipulated without changing the structure of the formulation that yields strong LP-relaxations and short solution times.

5.1 Metrics

In Chapter 3, we stated that an improved initial cut requires fewer aircraft, takes less time to generate, and better captures AMC's three operational objectives of *readiness*, *customer service*, and *net operating result (NOR)*. In Chapter 4, we illustrated that the CRM-C can be solved quickly, and we defined the composite variable generation procedures such that all composite variables satisfy the maximum port hold time restriction (thereby implicitly capturing customer service). Additionally, we use the metrics in *Table 5-1* to measure the characteristics of the solution output from CRM-C.

METRIC	DESCRIPTION	INTENDED TO MEASURE:
Aircraft	<i>The total number of each aircraft type used to operate the channel route network</i>	<i>Desire to use fewer aircraft</i>
System Ute Rate	<i>Average utilization rate of aircraft (with respect to the TWCF planning weight) over <u>all</u> flight legs in the planning horizon</i>	<i>NOR</i>
Operating Cost	<i>The total operating cost over all channel route missions</i>	<i>NOR</i>
AMC Hold Time	<i>The average AMC hold time over all cargo commodities</i>	<i>Customer Service</i>
Flying Hours	<i>The total number of flying hours for each aircraft type</i>	<i>Readiness</i>
TWCF Ute Rate	<i>Average utilization rate of aircraft (with respect to the TWCF planning weight) over all <u>CONUS</u> outbound flight legs in the planning horizon</i>	<i>Current AMC metric</i>
Missions	<i>The total number of channel route missions flown by each aircraft type</i>	<i>Current AMC metric</i>

Table 5-1: Metrics used to analyze CRM-C

OBJECTIVE (<i>minimize</i>)	DESCRIPTION
Aircraft (Acft)	<i>Minimize the number of aircraft that traverse the set of wrap around arcs (defined in §3.4.2)</i>
Operating Cost (Cost)	<i>Minimize the monthly aircraft operating costs of the channel route network</i>
Missions Flown (Msns)	<i>Minimize the total number of channel route missions flown over the planning horizon</i>
Unused aircraft capacity (Ute)	<i>Minimize the total unused aircraft capacity over the CONUS outbound flight legs (i.e., maximizing TWCF ute rate) over the planning horizon</i>

Table 5-2: Objective functions tested

As we discussed in Chapter 4, data for each composite variable are stored in a data structure. This allows the metrics in *Table 5-1* to be stored simply as fields within the data structure. Each metric is calculated during the composite variable generation, which allows the metric calculations to be modified – and even new metrics to be added – without changing the structure of CRM-C.

5.2 Objective Functions

One advantage of CRM-C is its ability to be solved quickly, which leads to two improvements to the current channel route planning process. First, it reduces the resources (i.e., man-hours) required to create an initial cut of the 30-day channel route schedule. Second, it provides the organic channel scheduler the ability to test the sensitivity of the channel route schedule to various objective functions (i.e., planning priorities) and route structures.

We evaluate four objective functions, which we describe in *Table 5-2*. Similar to the metrics in §5.1, the composite variable objective function coefficients for each of the objective functions in *Table 5-2* is stored as a field in the data structure created to store the composite variables.

One shortcoming of the current channel route planning process is the initial cut's propensity to change during the execution month. As described in Chapter 2, one way to address this weakness is to create an initial cut that requires fewer aircraft, thereby reducing the impact of aircraft availability on the channel route schedule. The objective function **Acft** in *Table 5-2* captures this idea. Another way to address this shortcoming is to minimize the total number of channel route missions flown, which is captured by objective function **Msns** in *Table*

5-2. Currently, channel route planners create an initial cut that will cover all cargo commodities and service all frequency channels, while maximizing the CONUS outbound utilization rate (i.e., TWCF ute rate). The end-of-month CONUS outbound utilization rates are typically near 1.0. We capture this process using the **Ute** objective function in *Table 5-2*. Finally, we use the **Cost** objective function as a proxy for NOR. *Table 5-3* summarizes the aircraft characteristics for each aircraft type.

Aircraft Type, f	Number of Aircraft Available, ρ^f	Operating Cost (\$/hr)	Capacity (tons)	TWCF Planning Weight (tons)
C5-DOV	12	19000	116	50
C17-CHS	15	9000	68	25
C141-WRI	4	8900	42	16
C17-RMS	3	9000	68	25
C130-RMS	15	6000	17	8

Table 5-3: Number of aircraft available of each aircraft type

5.3 Parameters

In addition to various objective functions, we also test different values of the parameters introduced in Chapter 4 and summarized in *Table 5-4*. It is important to note that the parameters in *Table 5-4* are inputs to the composite variable generation procedures developed in Chapter 4.

PARAMETER	ABBREVIATION	DESCRIPTION
Port Hold Time	<i>PHT</i>	<i>The maximum time that a cargo commodity is permitted to remain at its APOE after its Available to Load Time (ALT)</i>
Residual Aircraft Capacity Available	<i>Fract</i>	<i>The minimum amount of residual aircraft capacity that a Single-Route Composite mission must have available to be considered for inclusion in a Multi-Route Composite mission</i>
Flight Sequences	<i>Seq</i>	<i>The input set of Flight Sequences</i>
StartFreq	<i>SF</i>	<i>The start frequency of each generic mission</i>

Table 5-4: Summary of parameters tested

Changing the parameters simply changes the number of composite variables generated and the data (e.g., the constraint coefficients) associated with the composite variables. Consequently, the structure of CRM-C is preserved.

5.4 Computational Performance

The purpose of this section is to investigate the computational performance (i.e., the strength of the LP-relaxation and solution times) of CRM-C when the different objective functions are used, and to evaluate the influence that each parameter (which are used in the composite variable generation procedures described in Chapter 4) has on the size of CRM-C. We examine the influence of each parameter independent of the other parameters. All tests were run using a Pentium III 550 MHZ processor, 512 megabytes of RAM, and the commercial optimization software XPress-MP (version 12).

5.4.1 Objective Functions

We tested several data sets, each with a different number of variables. The objective function used to solve CRM-C impacts the strength of the LP-relaxation and the time required to obtain an integer solution, as shown in *Table 5-5*. Each column in *Table 5-5* corresponds to a different data set, which is created by changing a single parameter value in the *Baseline* data set. The *Baseline* data set has the following characteristics: $PHT = 2$, $Fract = 0.2$, $SF = 1$, and a set of flight sequences (*Seq Base*) taken from an existing real-world channel route schedule (which is the same set of flight sequences used in Chapter 3). The column heading for each data set in *Table 5-5* describes how the data set differs from *Baseline*. For instance, the data set *PHT1* has the PHT parameter set to 1 and all other parameters are the same as in *Baseline*. The data sets *SeqB* and *SeqE* consist of different route structures (e.g., direct routes, hub-and-spoke).

Minimizing the operating cost (**Cost**) yields the tightest LP-relaxations, and thus shortest solution times for all instances. Minimizing the number of aircraft (**Acft**), on the other hand, yields the weakest LP-relaxations and the longest solve times. It is interesting to note that even when **Acft** has a relatively tight LP-relaxation, the computational effort required in the branch-and-bound tree is significant (see data set *SeqE* for the **Acft** objective function in *Table 5-5*). It appears that the **Acft** objective function causes degeneracy, which occurs because there is little differentiation in the value of each column. In an attempt to address the degeneracy, we perturbed the objective function coefficient of each column slightly, which is how

<i>Data Set</i>		<i>Baseline</i>	<i>PHT1</i>	<i>Fract10</i>	<i>Fract05</i>	<i>SeqB</i>	<i>SeqE</i>	<i>SF3</i>	<i>SF4</i>
<i>Size</i>	Rows	1301	1301	1301	1301	1301	1301	1301	1301
	Columns	16982	11307	19691	22657	73106	81993	86735	142767
	Nonzeroes	197596	94722	256918	325018	1816220	2382102	1258700	2137261
<i>Obje- ctive func- tion</i>	<i>Acft</i>	LP (sec)	317	76	354	365	1528	1146	508
		*IP (sec)	1596	306	2386	3827	24058	29387	24446
		LP-IP Gap	0.046	0.021	0.087	0.132	0.197	0.031	0.139
	<i>Cost</i>	LP (sec)	154	61	187	204	1583	987	269
		*IP (sec)	116	17	107	146	2023	1030	720
		LP-IP Gap	0.022	0.010	0.024	0.017	0.031	0.0314	0.018
	<i>Msns</i>	LP (sec)	202	71	201	245	2078	955	401
		*IP (sec)	260	57	258	356	5768	3599	2508
		LP-IP Gap	0.028	0.015	0.034	0.027	0.067	0.059	0.033
<i>Ute</i>		LP (sec)	255	84	205	393	1857	1052	447
		*IP (sec)	322	129	337	443	5636	10241	4266
		LP-IP Gap	0.031	0.030	0.030	0.040	0.059	0.033	0.079

*All instances stopped after the first integer solution

Table 5-5: Influence of objective functions on the computational performance of CRM-C

we solved the **Acft** formulation reported in Table 5-5. Using **Msns** and **Ute** has varied results, but these objectives always perform computationally better than **Acft** and worse than **Cost**.

It should also be noted that we terminated each run after the first integer solution was found. The results in the remainder of this chapter are also based on the first integer solution. From a theoretical standpoint, we justify using the first integer solution, rather than the optimal integer solution, because of the tight LP-IP gaps illustrated in Table 5-5. From a practical standpoint, we justify using the first integer solution because any mathematical model aimed at aiding the channel route planners should yield solutions in a short amount of time. This is, in fact, a strength of CRM-C. Namely, CRM-C yields excellent solutions in a small amount of time. It should also be noted, however, that CRM-C allows the user to trade solution time for the desired degree of optimality. This is particularly important for the few instances that have large LP-IP gaps (see the data set *SF3* in Table 5-5).

5.4.2 Parameters

The purpose of this section is to examine the influence of each parameter on the size of CRM-C. We change the maximum allowed *port hold time* (*PHT*), *Fract* (i.e., amount of residual capacity that must be available for a single-route composite (SRC) mission to be considered for

inclusion in a multi-route composite (MRC) mission), the input set of *flight sequences*, and *StartFreq* (i.e., the number of time periods modeled per day). We present the effect of changing each parameter, independent of the other parameters in the data set *Baseline*. This leads to an overall picture of how the size of CRM-C fluctuates as the input parameters are changed, which we summarize in the last part of this section.

5.4.2.1 Port Hold Time

Decreasing the PHT (i.e., increasing the guaranteed customer service level) significantly reduces the set of SRC missions and MRC missions that are included in CRM-C, as *Figure 5-1a* illustrates. The PHT is enforced in the composite variable generation procedures and a more restrictive (i.e., lower) value decreases the number of feasible columns that must be considered in CRM-C.

5.4.2.2 Residual Aircraft Capacity Available

The amount of residual aircraft capacity that an SRC mission must have available to be considered for inclusion in an MRC mission does not significantly impact the size of CRM-C, as illustrated in *Figure 5-1b*. *Fract* defines the amount of residual aircraft capacity that an aircraft must have available to be considered for inclusion in an MRC mission. Specifically, *Fract* corresponds to the fraction of a cargo commodity's tonnage requirement. For instance, consider a cargo commodity with a tonnage requirement of 150 tons and let *Fract* = 0.10 (i.e., 10%).




Figure 5-1: Impact of (a) PHT and (b) Fract on the number of columns in CRM-C

This cargo commodity cannot be covered by an SRC mission because the tonnage requirement is greater than any aircraft's capacity. Consequently, an MRC mission is created to cover the cargo commodity. To construct the MRC mission, we piece together enough residual aircraft capacity from numerous SRC missions to cover the cargo commodity. We consider only those SRC missions that have at least 15 tons of residual capacity available. The *Fract* parameter captures the fact that, operationally, it is not desirable to split cargo commodities into small shipments. Smaller values of *Fract* increase the number of SRC missions that are candidates for MRC missions, which slightly increases the number of MRC mission columns in CRM-C, as *Figure 5-1b* illustrates.

5.4.2.3 Flight Sequences

The input set of flight sequences (from which generic missions and composite missions are created) plays a major role in determining the size of CRM-C. *Figure 5-2* illustrates the number of columns (all instances have 1301 rows) associated with three different sets of flight sequences. The *Seq Base* set of flight sequences contains a set of flight sequences taken from a real-world operational channel route schedule (the same set of flight sequences used in Chapter 4), which are mostly direct aircraft routes (i.e., between the APOE and APOD of a cargo commodity). The *Seq B* set of flight sequences contains all the flight sequences in *Seq Base* plus *regional-tour* flight sequences. A regional-tour flight sequence contains aerial ports that are located in a common geographic region (e.g., the Middle-East region, the Mediterranean region). Most of the regional-tour flight sequences contain a hub aerial port. *Seq E* consists solely of flight sequences that visit a hub aerial port, including the regional tour flight sequences from *Seq B*, that satisfy this restriction. *Example 5-1* illustrates the nature of the flight sequences in *Seq Base*, *Seq B*, and *Seq E*.

Example 5-1: Consider a channel route network depicted in *Figure 5-2*, which consists of six aerial ports. Aerial port B is a hub aerial port, aerial ports C and D are in the same geographic region, and aerial ports E and F are in the same geographic region. *Figure 5-2a* depicts the type of flight sequences in *Seq Base*. *Seq B* consists of the flight sequences in *Seq Base* (which are the light gray arrows in *Figure 5-2b*), plus the flight sequences depicted by the black arrows in *Figure 5-2b*. *Seq E* consists of the type of flight sequences depicted in *Figure 5-2b* (which are the light gray arrows in *Figure 5-2c*), plus the flight sequences illustrated by the black arrows in *Figure 5-2c*.




Figure 5-2: Structure of flight sequences in Seq Base, Seq B, and Seq E

Although we do not allow cargo to transfer between aircraft at hub aerial ports, we do allow cargo originating at the hub to be picked-up by any aircraft visiting the hub. In general, the flight sequences added to *Seq B* and in *Seq E* contain more OCONUS aerial port stops per takeoff from a CONUS aerial port. Although cargo is not permitted to transfer aircraft at hub locations, the flight sequences in *Seq B* and *Seq E* are reflective of a hub-and-spoke route structure.

The *type* of flight sequences has more of an impact on the number of columns in CRM-C than the *number* of flight sequences, as *Figure 5-3* illustrates. Specifically, flight sequences that contain hub aerial ports increase the size of CRM-C because the hub aerial ports are also *high cargo generation aerial ports*. A high cargo generation aerial port is the destination of a cargo commodity whose tonnage requirement cannot be completely covered by any single aircraft. Therefore, more flight sequences that contain a high cargo generation aerial port increase the number of SRC missions that are candidates for inclusion in an MRC mission variable (i.e., there are more combinations of SRC mission variables that will cover the cargo commodities that have tonnage requirements exceeding the capacity of all aircraft).




Figure 5-3: Impact of flight sequences on the number of columns in CRM-C

5.4.2.4 *StartFreq*

The *StartFreq* (SF) parameter plays a key role in determining the amount of effort required to create an operationally feasible schedule from the output of CRM-C (see §4.5). *StartFreq* represents the number of times per day that each generic mission can be initiated. For instance, consider the generic mission with the flight sequence $A \rightarrow B \rightarrow C \rightarrow D \rightarrow A$. When *StartFreq* = 1, we generate only one SRC mission variable, which has a start time of 0000 hours. When *StartFreq* = 3, we generate three SRC mission variables, which have start times 0000 hours, 0800 hours, and 1600 hours. Clearly, larger values of *StartFreq* correspond to a more detailed representation of the real-world system. However, larger values of *StartFreq* also increase the size of CRM-C, as *Figure 5-4* illustrates.

There is a balance between fidelity and tractability. Increasing the value of *StartFreq* does not impact the solution quality (as measured by the metrics introduced in §5.1), which we illustrate in §5.5. However, it is desirable to obtain a solution that can be easily transitioned into an implementable schedule. Therefore, solving CRM-C with *StartFreq* = 1 at the beginning of each month will successfully determine the route structure of the initial cut, and determine the number of required aircraft, AMC hold time, aircraft utilization rates, and operating costs. To obtain a solution that can be easily transitioned in an operationally feasible schedule, CRM-C must be solved with larger values of *StartFreq*. However, larger values of *StartFreq* cause CRM-C to be intractable for the **Acft** objective (see *Table 5-5*). To overcome this difficulty, CRM-C could be solved over a smaller planning horizon (e.g., a weekly time horizon) with larger values




Figure 5-4: Impact of *StartFreq* on the number of columns in CRM-C

of *StartFreq* and the set of flight sequences used in the solution to the monthly problem (where *StartFreq* = 1). This would yield a tractable formulation with a high level of fidelity. We discuss this type of *hierarchical decomposition* in Chapter 6.

5.4.2.5 The Effect of Parameter Values on the Size of CRM-C

Table 5-6 summarizes the impact of each parameter on the size of CRM-C, where ↑ represents an increase in value, ↓ represents a decrease in value, and ↔ represents that the value is unchanged.

5.4.2.6 Simultaneous Parameter Changes

The advantage of simultaneously varying parameters is improved solution quality. However, as Figure 5-5 illustrates, there are combinations of parameters that yield large – and probably intractable – instances of CRM-C. It would not be necessary to attempt to solve a large instance of CRM-C if it is not expected to improve the solution quality. The remaining sections in this chapter are aimed at determining the parameters that yield improvements in solution quality.

PARAMETER	DIRECTION OF CHANGE IN PARAMETER VALUE	INFLUENCE ON CRM-C	
		Number of SRC variables	Number of MRC variables
<i>PHT</i>	↑	↑	↑
	↓	↓	↓
<i>Fract</i>	↑	↔	↓
	↓	↔	↑
<i>Flt Sqncs</i>	number of flight sequences ↑	↑	depends on the type of flight sequences added
	number of flight sequences through a hub aerial port ↑	↑	↑
<i>SF</i>	↑	↑	↑
	↓	↓	↓

Table 5-6: Summary of parameters' effect on the size of CRM-C




Figure 5-5: Impact of Multiple Parameters on the number of columns in CRM-C

5.5 Solution Analysis

The purpose of this section is to examine the impact that the different objective functions and the different parameter values have on the characteristics of the solution to CRM-C, as measured by the metrics introduced in §5.1. This section also highlights the wide range of analysis that can be performed using CRM-C, with only slight changes to the inputs of the composite variable generation procedures.

5.5.1 Objective Functions

The objective function used has a significant impact on the solution, as *Figure 5-6* illustrates. The **Acft** objective yields a solution that uses mostly large-capacity aircraft (i.e., the C5 and the C17), which fly most of the missions and accumulate most of the flying hours. The **Msns** objective yields a solution consisting of aircraft of all capacities; however, the large-capacity aircraft fly most of the missions and accumulate most of the flying hours (depicted in *Figure 5-6b* and *Figure 5-6c*).

Using the **Cost** objective function, the mix of aircraft in the solution is similar to the mix of aircraft in the solution obtained using the **Msns** objective function. In the solution using the **Cost** objective function, however, the large-capacity aircraft – which are more expensive – fly

fewer missions. Specifically, more C141-WRI flight hours are used instead of C5-DOV flight hours and more C130-RMS hours are used instead of C17-RMS flight hours. This tradeoff between these aircraft is a property of the data. In general, the tonnage requirements of the cargo commodities can be carried by the smaller capacity aircraft, which have lower operating costs (see *Table 5-3*).

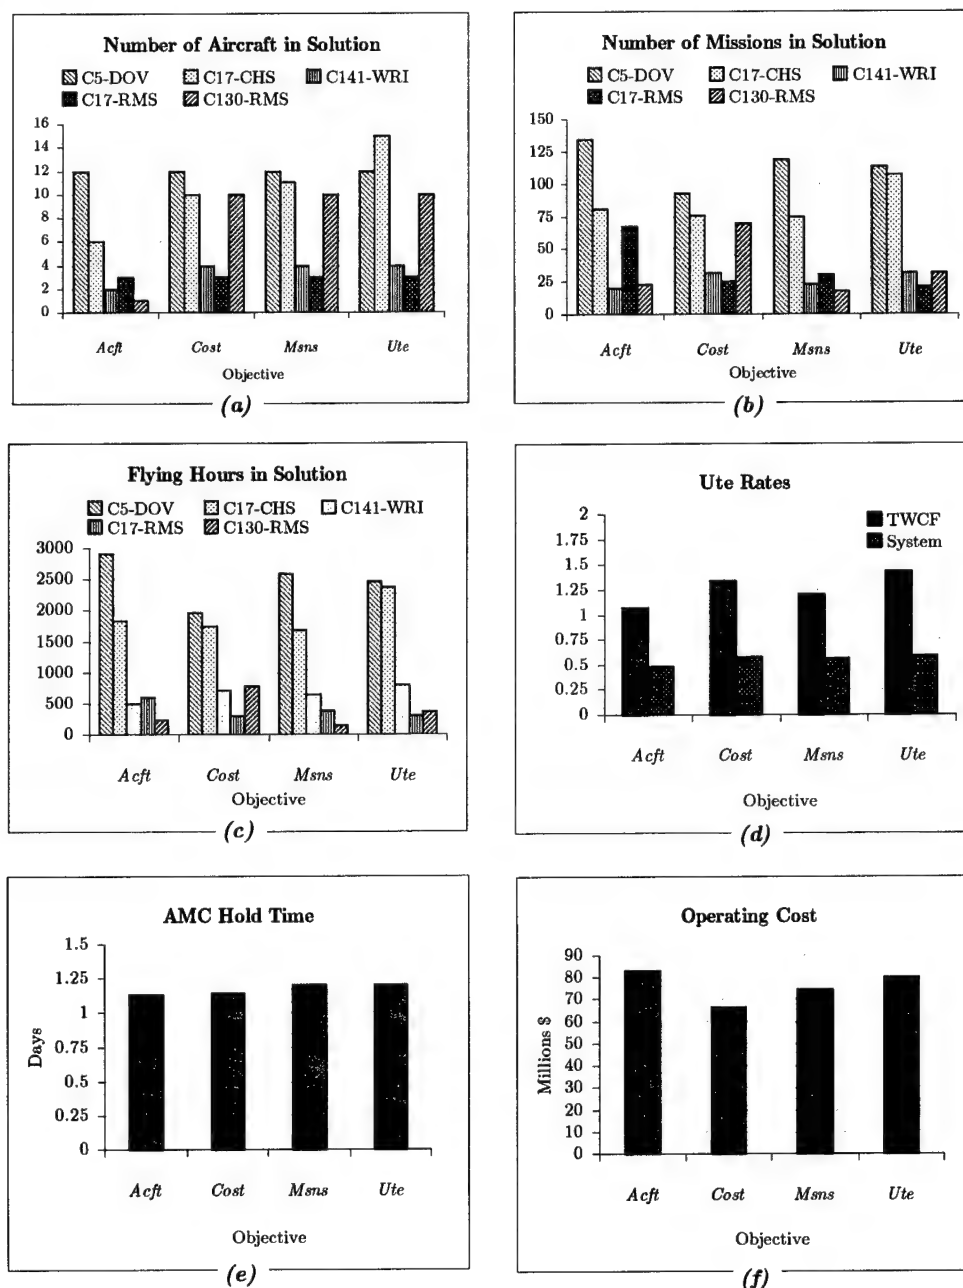


Figure 5-6: Metric values vs. Objective Function

The **Ute** objective – which minimizes the unused aircraft capacity on the CONUS outbound flight leg – yields a solution that uses all available aircraft, and almost all missions are operated by large capacity aircraft that are stationed at CONUS aerial ports.

In general, higher utilization rates (see *Figure 5-6d*) are associated with longer AMC hold times (see *Figure 5-6e*). The **Cost** objective function achieves the lowest overall operating cost by using fewer total flying hours with relatively few missions, which have high utilization rates. Although the **Acft** objective function yields a solution with fewer aircraft than the other objective functions, the solution relies on the C5 – which is the most expensive aircraft – to fly nearly half the total flying hours. As a result, the operating cost associated with the **Acft** objective is the largest (see *Figure 5-6f*).

Table 5-7 summarizes the performance of each objective function. For each metric, the four objective functions are ranked, where a rank of 1 corresponds to the objective function that performed the best with respect to the metric and a rank of 4 corresponds to the objective function that performed the worst. We consider fewer aircraft, higher utilization rates, more flying hours, and lower cost desirable. It is not clear whether a solution with more channel route missions is better, so we define a rank of 1 for the *Missions* metric to represent the solution with the largest number of channel route missions flown.

Table 5-7 highlights an interesting relationship between AMCs three operational objectives (readiness, customer service, NOR). For each objective function, the model performs well with respect to at most two of the operational objectives (represented by the bold, italic values in the grayed fields in *Table 5-7*). Specifically, the relationships are:

1. High levels of customer service and readiness can be achieved, but at a high cost.
2. High levels of customer service can be achieved at a low cost, but readiness will be low.

Objective	Number of Acft	Current AMC Metrics		Readiness	NOR		Customer Service
		Missions	TWCF Ute Rate	Flying Hours	System Ute Rate	Cost	AMC Hold Time
Acft	1	1	4	2	4	4	1*
Cost	3*	3	2	3	1*	1	1*
Msns	3*	2	3	4	3	2	2
Ute	4	1	1	1	1*	3	2

*Tied

Table 5-7: Summary of objective function performance

Note that using the **Ute** objective function performs the best with respect to the metrics used in the current channel route planning process. The major drawback of the **Ute** objective function is that it yields solutions that use all available aircraft.


5.5.2 Parameters

In the following sections, we investigate the influence that each parameter – PHT, *Fract*, flight sequences, *StartFreq* – has on the values of the metrics (i.e., the solution characteristics) introduced in §5.3. Initially, we analyze the influence of each parameter independent of the other parameters. Specifically, we use the *Baseline* data set and vary each parameter while holding the values of the other parameters constant. This leads to an overall understanding of the behavior of the metrics as the objective function and input parameters are changed. In the final part of this section, we analyze the impact of simultaneously varying several parameters.


5.5.2.1 Port Hold Time

Decreasing the maximum allowed PHT means that cargo will spend less time at the APOE's. This requires aircraft to move the cargo quickly, reducing the amount of time that cargo can build-up at an aerial port. Because the tonnage requirements of the cargo commodities are generally much smaller than the aircraft capacities, a shorter PHT will result in lower aircraft utilization. The end result is that more aircraft are required (see *Figure 5-7a*) and more missions must be flown (see *Figure 5-7b*) to keep the cargo moving through the system quickly. This reduces the total time that cargo spends in the system (see *Figure 5-7e*), but the operating cost increases (see *Figure 5-7f*). As expected, decreasing the PHT results in higher customer service (i.e., shorter AMC hold Times) – see *Figure 5-7e* – but at the expense of more aircraft, more missions, more flying hours, higher operating costs, and lower utilization rates.


Figure 5-7 also illustrates that a second order objective might need to be identified. For instance, assume that we want the solution with the fewest flying hours. *Figure 5-7c* illustrates that the solution obtained using the Cost and Msns objective functions have the same number of flying hours. Using the other metrics, however, the solution obtained using the **Cost** objective function clearly dominates the solution obtained using the **Msns** objective function. Specifically, the solution obtained using the **Cost** objective function requires fewer aircraft (see *Figure 5-7a*), results in higher aircraft utilization rates (see *Figure 5-7d*), results in a shorter AMC Hold Time (see *Figure 5-7e*), and has a lower operating cost (see *Figure 5-7f*).




(a)




(b)




(c)



(d)



(e)



(f)

Figure 5-7: Metric values vs. Port Hold Time (PHT)




Figure 5-8: Metric values vs. Fraction of Available Aircraft Capacity

5.5.2.2 Residual Aircraft Capacity

Figure 5-8 illustrates that decreasing the value of *Fract* does not significantly improve the solution. Decreasing the value of *Fract* (i.e., creating more feasible MRC missions) slightly decreases the cost (see Figure 5-8c), but the same number of aircraft (see Figure 5-8a) and missions (see Figure 5-8b) are required. Additionally, the aircraft utilization rates increase slightly (see Figure 5-8a).

5.5.2.3 Flight Sequences

The set of flight sequences (described in §5.4.2.1) significantly impacts the solution to CRM-C. Flight sequences that resemble a hub-and-spoke structure (i.e., *Seq B* and *Seq E*) drastically reduce the number of aircraft, missions, and flying hours (see *Figure 5-9*). The exception is when we use the **Acft** objective function. In this case, the number of aircraft required increases slightly (see *Figure 5-9a*) because the longer flight sequences prevent the aircraft from returning to the CONUS aerial ports as quickly as when flying more direct flight sequences (i.e., the *Seq Base* flight sequences).




Figure 5-9: Metric values vs. Flight Sequences

Figure 5-9 illustrates that flight sequences more like the hub-and-spoke route structures of *Seq B* and *Seq E* increase the utilization rates and decreases the operating cost. However, the cargo remains in the system longer. In general, it appears that *Seq B* and *Seq E* yield superior results because these route structures require fewer resources, allow higher aircraft utilization, and have lower operating costs. However, AMC must maintain a minimum level of readiness – which is achieved through flying hours – and hub-and-spoke route structures require fewer flying hours.

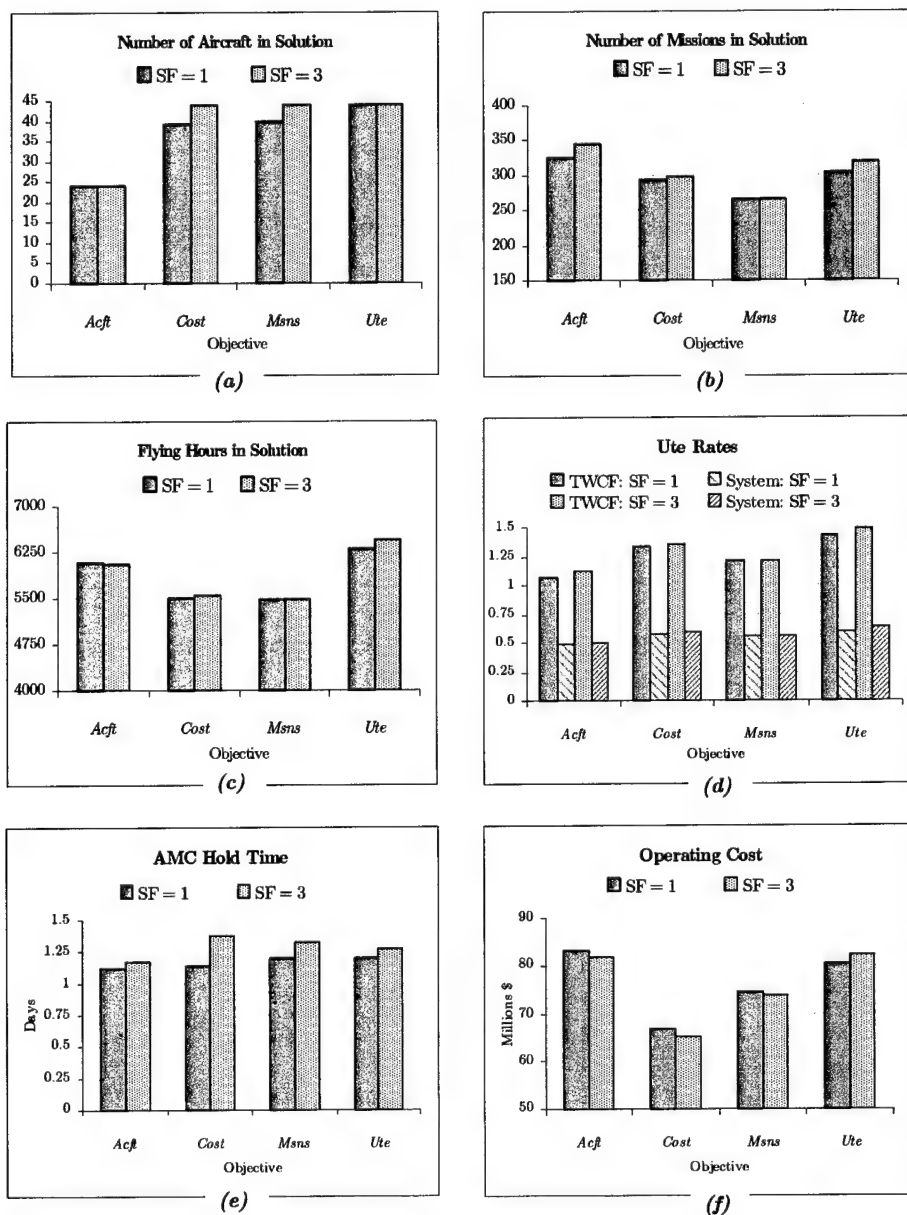


Figure 5-10: Metric values vs. StartFreq

5.5.2.4 StartFreq

The major trend apparent in *Figure 5-10* is that the value of *StartFreq* does not significantly influence the characteristics of the solution to CRM-C. There are small changes in the solution characteristics, which include an increase in AMC Hold Time (see *Figure 5-10e*), slightly higher utilization rates (see *Figure 5-10d*), and slight changes in operating costs (see *Figure 5-10f*). The *StartFreq* parameter has an impact on the effort required to transition the output from CRM-C into an operationally feasible schedule. Increasing *StartFreq* captures more detail, but at the expense of increased formulation size. As *Figure 5-10* illustrates, the benefit of increasing the fidelity is not rooted in increased solution performance. Rather, increasing *StartFreq* is desirable only because it reduces the effort required to transition the output of CRM-C into an operational schedule.

5.5.2.5 The Effect of Parameter Values on the Solution of CRM-C


Table 5-8 summarizes the influence that the input parameters have on the metrics (where ↑ represents an increase, ↓ represents a decrease, and ↔ represents unchanged) as well as the impact on the size of the formulation from *Table 5-6*.

PARAMETER	CHANGE IN PARAMETER VALUE	CHANGE IN METRIC VALUE						MODEL SIZE
		<i>Aircraft</i>	<i>Missions</i>	<i>Flying Hours</i>	<i>Ute Rates</i>	<i>Hold Time</i>	<i>Cost</i>	
<i>PHT</i>	↓	↑	↑	↑	↓	↓	↑	↓
<i>Fract</i>	↓	↔	↔	↔	Slight ↑	Slight ↑	Slight ↓	↑
<i>Flt Sqncs</i>	<i>Direct</i>	↑	↑	↑	↓	↓	↑	↓
	<i>Hub-and-Spoke</i>	↓	↓	↓	↑	↑	↓	↑
<i>SF</i>	↑	↔	↔	↔	Slight ↑	Slight ↑	Slight ↓	↑


Table 5-8: Summary of parameters' effect on the solution of CRM-C

5.5.2.6 Simultaneous Parameter Changes


Table 5-8 illustrates the effect of each parameter on the solution to CRM-C and the size of the formulation. The parameters can be changed simultaneously in many combinations. For instance, if the objective (i.e., the channel route planner's priority) is to minimize cost, *Table 5-8* illustrates that decreasing the *Fract* parameter and using a hub-and-spoke route structure both decrease the operating cost. However, both increase the size of the formulation. We test this combination, using *Fract* = 0.10 and *Seq B* (and other *Baseline* parameter values) and the




(a)




(b)




(c)



(d)



(e)



(f)

Figure 5-11: Effect of simultaneously varying parameters on metrics

Cost objective function. The resulting formulation's size and solution times are summarized in Table 5-9.

<i>Data Set</i>		<i>Baseline</i>	<i>Seq B</i>	<i>Fract10</i>	<i>Seq B, Fract10</i>
Size	Rows	1301	1301	1301	1301
	Columns	16982	73106	19691	101637
	Nonzeroes	197596	1816220	256918	3091554
Cost	LP (sec)	154	1583	187	1939
	*IP (sec)	116	2023	107	3139
	LP-IP Gap	0.022	0.031	0.024	0.0362

**Stopped after the first integer solution*

Table 5-9: Size and solution time with *Fract* = 0.10, *Seq B*, and *Cost* objective

Decreasing the *Fract* parameter and using the *Seq B* route structure yields slight improvements compared to changing each of the parameters individually and compared to the *Baseline* values of the parameters. Specifically, the number of aircraft, missions, flying hours all decrease, as Figure 5-11 illustrates, but at the expense of increased computational time. Furthermore, the utilization rates increase slightly (see Figure 5-11a), resulting in slightly longer AMC Hold Time (see Figure 5-11a). Additionally, the operating cost decreases because we've minimized the operating cost over a larger set of variables.

5.6 Summary

In this chapter we evaluated the computational performance and the solution characteristics of CRM-C under different objective functions and parameter values. The results presented in this chapter illustrate that CRM-C is useful for identifying and quantifying tradeoffs made between different planning strategies (i.e., objective functions) with respect to a set of metrics that reflect AMC's three operational objectives.

In general, the objective function used to solve CRM-C affects the strength of the LP-relaxation and the characteristics of the solution. Specifically, the input set of flight sequences plays a major role in defining the size of the resulting formulation and characteristics of the solution. The *StartFreq* parameter also plays a major role in the size of CRM-C. It does not, however, significantly change the solution quality. Increasing the value of *StartFreq* (i.e., increasing the model fidelity) yields a solution that requires less effort to transition into an

implementable channel route schedule, but at the expense of increased model size and increased computational effort.

This chapter also highlights a major advantage of composite variable modeling – composite variables can be modified through the composite variable generation procedures without destroying the structure of the model. Consequently, we are able to change the problem with minimal effort (e.g., modifying the fidelity, changing the objective function, testing new route structures) while maintaining the formulation's structure that results in strong LP-relaxations and short solution times.

[This Page Intentionally Left Blank]

6 Towards a Decision Support Tool

In Chapter 5, we illustrated CRM-C's ability to handle different planning objectives. Furthermore, we demonstrated that simply modifying the composite variable generation procedures allows us to address various aspects of channel route network design. To this end, we've successfully illustrated that an optimization-based planning approach can aid channel route planners in designing the channel route network.

The purpose of this chapter is to highlight how this research could be transitioned into an *optimization-based decision support tool*. In our discussion, we outline avenues of future research related to enhancing the work developed in this thesis. This chapter is organized into four sections. First, we define our concept of an optimization-based decision support tool. In the second section, we discuss the integration of an optimization-based decision support tool with existing software tools that are used by channel route planners, and we highlight how such tools would fit into the current planning process. Third, we describe the range of specific capabilities that an optimization-based decision support tool would be well-suited to handle. In the final section, we discuss specific enhancements that would allow CRM-C to capture more real-world attributes in the channel route network and that would allow CRM-C to perform the desired capabilities that we highlight.

6.1 What is a Decision Support System

Our concept of an *optimization-based decision support tool* is best defined by the characteristics that such a tool should possess. An optimization-based decision support tool is built around mathematical models that capture the important features of a real-world system and assist in the process of making decisions pertaining to the system design and operation. It is important to note that an optimization-based decision support tool does not necessarily require that the system be designed, or operated, at the "optimal" level. In fact, the tool should support rather than automate the decision process. This implies user interaction. Furthermore, an optimization-based decision support tool should respond to changes both in the operations of the day-to-day system and to changes in organizational planning processes.

6.2 Integrating a Decision Support Tool within the Current Planning Process

In §2.3.2, we outlined the current software tool used by channel route planners, which is called the Consolidated Air Mobility Planning System (CAMPS). CAMPS is an interactive environment that allows channel route planners to create and modify channel route missions. CAMPS significantly reduces the laborious tasks that an organic channel route scheduler faces, including gathering and consolidating all the information required to create an initial cut. CAMPS does not, however, currently use an automated scheduling algorithm to help the organic channel route scheduler determine the best set of aircraft routes and cargo flows. One of CAMPS strengths – which an optimization-based decision support tool should utilize – is that it interfaces with several databases that house information about aircraft characteristics, crew rules, and the schedules from other mission areas. The databases with which CAMPS interfaces contain the data required to build composite variables. Integrating an optimization-based decision support tool within the existing software environment would prevent the need for a completely new software package. An optimization-based decision support tool could, in fact, be as simple as an additional menu item in the existing software environment.

It is also important to note that an optimization-based decision support tool would not disrupt the flow of the current planning process, which is illustrated in *Figure 6-1a*. In other words, an optimization-based method would not require changes to the tasks that each channel route planner performs in generating the initial cut. It would simply change the order in which the CONUS cargo bookie and the barrelmaster provide inputs to the organic channel scheduler.

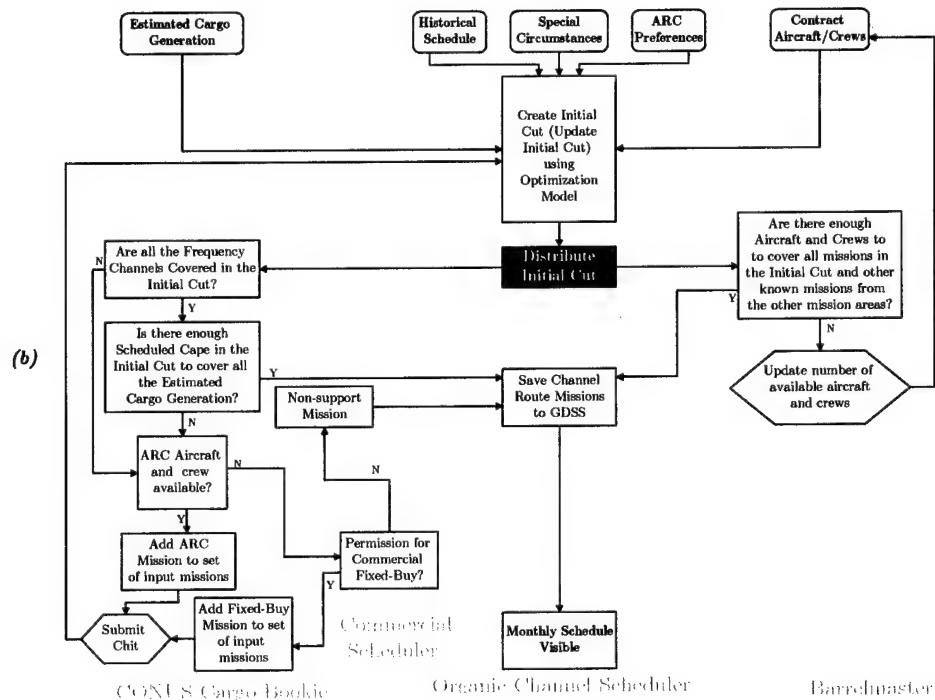
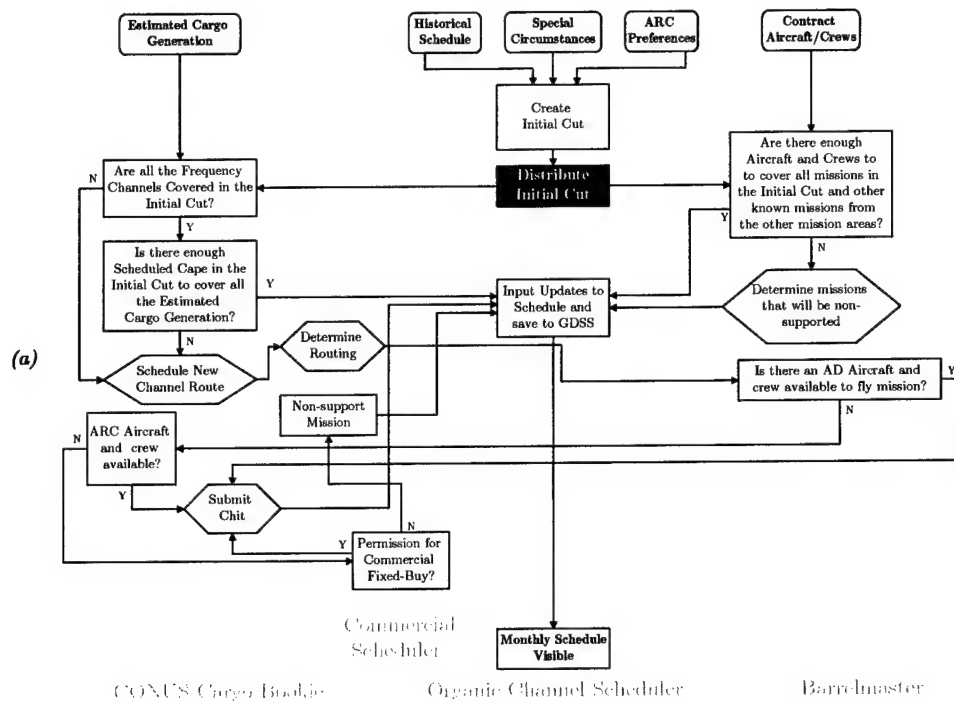


Figure 6-1: Planning (a) without and (b) with an optimization-based decision support tool

An optimization-based approach allows all of the inputs to be considered simultaneously rather than sequentially. In the current channel route planning process (illustrated in *Figure 6-1a*) several iterations between the organic channel route scheduler, the CONUS cargo bookie, and the barrelmaster are required before all the relevant inputs are considered. In contrast, when an optimization-based approach is used (see *Figure 6-1b*), the first iteration of the initial cut would account for all the necessary inputs. This reduces the time required to generate an initial cut.

It is important to note, however, that the work in this thesis was not specifically designed to dovetail with the current scheduling process. The work in this thesis was built to the functional requirements (i.e., the inputs and outputs) of the channel route planning process. As a result, an optimization-based decision support tool with this research at its core has the ability to change and adapt to new business processes and organizational structure.

6.3 Desired Capabilities

Using an optimization-based decision support tool would allow a wide range of analysis to be conducted. The purpose of this section is to outline several planning functions in which an optimization-based decision support tool would aid channel route planners. We organize our discussion to correspond to three levels of decision making. The levels are defined by the frequency with which decisions are made. First, we discuss capabilities related to the generation of the initial cut, which are decisions made once per month. Second, we look at capabilities that would aid decision making during the execution month, which are decisions made every few days to once a week. Finally, we discuss capabilities that would aid longer-term decisions, which are made quarterly, seasonally, or annually. Each of these areas is a suggested avenue of future research.

6.3.1 Initial Cut Decision Support

Each month, an initial cut is generated (discussed in Chapter 2). At this level, a decision support tool should allow an organic channel route planner to address issues such as:

Aerial port closures or decreased throughput capabilities – channel route planners should be able to exclude any aerial ports or specify when they are available (e.g., long-term runway construction might close an aerial port for several weeks).

Varying levels of aircraft availability – a decision support tool should be able to capture the fact the aircraft availability might vary from week to week (e.g., a major exercise consumes a number of aircraft, but only during the third week of the month).

Changes in overfly restrictions – at times, countries change their overfly rules. This changes the flight path that aircraft must fly between two aerial ports, which could be accounted for by either increasing or decreasing the flight times on the effected flight legs.

Each of the above situations can be handled within the cargo generation procedures and would not change the structure of CRM-C.

6.3.2 Execution Month Decision Support

A future avenue of research is to explore how composite variable modeling can be used to make decisions within the execution month as the state of the system changes. A decision support tool would be extremely helpful in these situations because the system-wide impact of various solutions could be quickly quantified and compared. Formulating a model to be used during the execution month is challenging because one of the objectives of the problem is to minimize disruption to the existing channel route schedule. Planning with a *rolling horizon* is technique used to model systems that are continuously changing. In this method, each time re-planning occurs, a new plan is generated and the old plan is spliced into the new plan such that the transition from the old new plan to the new plan is seamless. In the channel route network, re-planning would be triggered either by an event that changes the state of the system or by a change in estimates of future states. We discuss events below that might trigger re-planning.

Aerial Port Closures – during the execution month, unknowns such as weather can close an aerial port with little forewarning. A decision support tool should allow a channel route planner to analyze the impact of such an event. Specifically, the decision support tool should identify the effect on the system performance (e.g., AMC hold time, utilization rates) and provide insights on how to re-route cargo to ensure customers at other aerial ports are serviced.

Aircraft Breakdown – military aircraft that experience mechanical failures can sit at an aerial port for days waiting for a maintenance crew and spare parts to be flown to the location.

Opportune Missions – occasionally, aircraft from other mission areas are available to move cargo in the channel route system. A decision support tool should be able to identify the effect this has on the channel route missions that are originally scheduled to transport the cargo.

Commercial Expansion Buy – there are periods when there is not enough military airlift to cover all the cargo and service all frequency requirements. In these situations, commercial carriers are contracted to transport cargo. A decision support tool should have the ability to test different flight sequences to minimize the cost of the expansion-buy mission. An optimization-based approach would be able to analyze the system-wide impact of additional aircraft capacity. For instance, in the current system, if a channel route mission is canceled as a result of an aircraft shortage, rather than contracting a commercial aircraft to fly the exact channel route mission canceled, it might be less expensive to re-route a military aircraft to cover the canceled mission and contract a commercial carrier to cover a different channel route mission.

Aircraft Availability – during the execution month, a channel route mission can be canceled because the aircraft scheduled to fly the mission must be used to fly a higher priority mission. In this situation, the lowest priority mission will be non-supported. There are often several channel route missions that are candidates to be non-supported. A decision support tool should be able to analyze the impact of canceling different missions.

Changes in cargo demand levels – formal methods are not currently used to forecast the demands that are used to create the initial cut. Often, the amount of cargo to be moved is less than expected or more than expected. This information is not known, however, until the cargo arrives at the aerial port. A decision support tool should be able to analyze the impact of fluctuations in the cargo demand levels.

6.3.3 Strategic Decision Support

Another area that can benefit from composite variable modeling is strategic decision making (i.e., decisions made quarterly or annually). An optimization-based decision support tool should be able to handle situations such as the following.

Route structure – Hub-and-spoke vs. Direct: The route structure has a major influence on flying hours. A decision support tool should be able to evaluate different

route structures. Furthermore, a decision support tool should have the capability to determine the optimal fleet mix at hub locations.

Commercial Fixed Buy Decision Support – on an annual basis, specific routes are offered to commercial carriers, which become the fixed-buy missions (see §2.2.3.2). A composite variable formulation is well suited to model this problem because the formulation can be solved quickly and the inputs can be manipulated easily, which would allow various fixed-buy route structures to be tested by simply modifying a single input file.

6.4 Enhancing CRM-C

Before the research in this thesis is used within a decision support tool, we suggest that several enhancements be explored that will improve the computational performance and executability of the solution.

6.4.1 Hierarchical Decomposition

In Chapter 5, we tested CRM-C using a small data set. As the size of the problem increases (i.e., additional aerial ports and cargo commodities are considered), a hierarchical decomposition of the channel route planning problem should be explored. This suggestion is based on the observation that increasing the fidelity of the model (i.e., increasing the value of the *StartFreq* parameter) is desired because it reduces the effort required to transition the output into an operational schedule. However, increasing the fidelity in this way also increases the model's size, rendering the model more difficult to solve. Increasing the fidelity does not, however, drastically affect the solution quality. This behavior can be exploited using hierarchical decomposition.

One way to hierarchically decompose the problem would be to solve the problem in two stages. In the first stage, a *monthly problem* is solved at a relatively low level of fidelity (i.e., *StartFreq*=1) to determine the general channel route structure. Solving the monthly problem would, for example, capture the interactions between the activities across weeks throughout a month (e.g., between the first and second weeks or between the first and fourth weeks). In the second stage, a *weekly sequencing problem* is solved at the beginning of each week. The weekly sequencing problem would have the same structure as CRM-C (i.e., the monthly problem), but the *StartFreq* parameter (see §4.4.2) would be increased to capture a higher level of detail. The

input set of flight sequences for the weekly sequencing problem would be limited to the flight sequences used during that week in the solution to the monthly problem.

6.4.2 Extreme Route Models and Column Generation

In addition to hierarchical decomposition techniques, we suggest that other techniques – such as the two we mention here – be explored that exploit the structure of CRM-C and allow larger, more detailed formulations to be solved.

Armacost [3] and Armacost et al. [4] transition from a traditional network design formulation to a composite variable formulation by first formulating an *extreme route formulation*. The extreme route formulation exploits the no-ramp transfer assumption. Specifically, an extreme route represents an extreme allocation of an aircraft's capacity to a subset of the cargo commodities. Armacost [3] and Armacost et al. [4] illustrate that feasible cargo flows are convex combinations of extreme aircraft routes, as long as ramp transfers are not permitted. Constraints in the extreme route formulation ensure that the combination of extreme routes is integral. Consequently, the integrality restriction on the decision variables can be relaxed. This technique should be explored as a possible method to deal with large channel route scheduling problem instances.

Column generation techniques play a key role in optimization problems that have a structure similar to CRM-C (i.e., a large number of columns and relatively few rows). In this thesis, we enumerated the composite variables, using the ideas of dominance and focusing our attention on sets of composite variables with specific properties. Even these techniques sometimes fail to restrict the set of feasible composites to a manageable size. Column generation techniques would allow larger, more detailed instances of the channel route scheduling problem to be solved.

6.4.3 Additional Composite Variables

In Chapter 4, we made the following assumptions, which exclude several aspects of the real-world channel route network.

Assumption 1: basic aircraft crews operate all channel route missions and remain with an aircraft for the duration of the channel route mission.

Assumption 2: Aerial refueling is possible, but not explicitly considered by CRM-C.

Assumption 3: Air National Guard and US Air Force Reserve (ARC) aircraft and aircrews are not considered.

Assumption 4: cargo cannot transfer between aircraft at aerial ports.

Relaxing these assumptions has the potential to improve the quality (i.e., reduce the number of aircraft and/or decrease the operating costs) of the solution to CRM-C, but requires that the composite variable generation procedures be modified to account for more complex composite variables. Furthermore, introducing more composite variables increases the size of the formulation. To deal with a large number of variables, column generation techniques for composite variables should be explored.

In the following sections, we discuss the types of composite variables that should be added to CRM-C. It is important to note that although the composite variables we discuss below are more complex than the composite variables that we consider in this thesis, they do not change the structure of the formulation.

6.4.3.1 Prepositioned & Augmented Crew Composites

In previous chapters, we assume that basic crews operate all channel route missions. Channel route planners, however, often take advantage of the fact that an augmented crew can extend the allowable crew duty day, as described in §2.2.4. Additionally, channel route planners take advantage of prepositioned crews, which allow an aircraft to move continuously through the channel route network (see §2.2.4).

In previous chapters, we also assume that crews remain with an aircraft for the duration of a channel route mission. As a result, we enforce the minimum crew rest requirement as constraint on the aircraft flows during the composite variable generation procedures. Allowing prepositioned crews would remove the crew rest requirement for the prepositioned crew composite variables. Additionally, a constraint would be added to CRM-C to limit the number of augmented crews that are used in an optimal solution. In a decision support environment, the user should be able to select the set of aerial ports where prepositioned crews are located.

6.4.3.2 Aerial Refueling Composites

Currently, channel route planners use experience to guide their decision about the channel route missions that can use aerial refueling (AR). The number of channel route missions that can use AR is restricted because the number of AR aircraft is limited, and aircraft flying in other mission areas (i.e., SAAMs, contingencies) must also utilize these resources. To plan a mission that uses AR, channel route planners use predetermined *AR tracts*. An AR tract

is simply an area – defined by a set of coordinates – where aircraft rendezvous with AR aircraft to refuel.

An easy way to incorporate AR into CRM-C is to increase the flight time for the flight legs on which aircraft use AR. For instance, consider aerial ports A and B and aircraft f_1 . If aircraft f_1 does not have the range to fly between A and B, the flight time can be increased to include the amount of time that it takes the aircraft f_1 to aerial refuel on the flight leg between aerial port A and aerial port B. A list would be maintained of the flight legs on which aircraft are able to use AR. Each AR composite variable is flagged and post processing indicates the channel route missions requiring AR. In a decision support environment, the user should be able to specify the aircraft and flight legs on which AR applies, as well as the maximum number of channel route missions that are permitted to use AR.

6.4.3.3 ARC Composites

ARC flying squadrons are not required to fly channel route missions. ARC squadrons are, however, a significant resource used by channel route planners. A decision support tool should incorporate ARC missions. Two methods could be used to model ARC missions. In the first method, ARC missions could be modeled similar to the way fixed-buy missions are modeled in Chapter 4. Using this method, an ARC mission with multiple maximal loadings would be input into CRM-C and constraints similar to (4.9) would be added to ensure that exactly one of the maximal loadings is in the solution. Because ARC flying squadrons are not required to fly channel route missions, channel route planners often attempt to accommodate ARC flying squadrons by honoring their channel route mission preferences. This method would allow channel route schedulers to easily account for ARC preferences. In the second method, ARC aircraft would be added to the list of input aircraft and the model would determine the routes the ARC aircraft are assigned to fly.

6.4.3.4 Ramp Transfer Composites

Currently, cargo is not permitted to transfer between aircraft at aerial ports, which is referred to as a *ramp transfer*. Armacost [3] and Armacost et al. [4] illustrate that ramp transfers yield significant cost savings for a major US carriers' next-day air network. Allowing ramp transfers increases the size of the problem by introducing more variables, and increases the complexity of the composite variable generation procedures. Allowing ramp transfers also assumes that aerial ports have sufficient infrastructure required to transfer cargo between aircraft. In a decision support environment, the user should be permitted to select the set of aerial ports where ramp transfers can occur. Furthermore, the user should be able to specify the

amount of cargo that can be transferred between aircraft (e.g., no more than fifty percent of cargo on an aircraft can be transferred). Ramp transfers increase the minimum required ground time, which would easily be accounted for in the composite variable generation procedures.

6.5 Summary

This chapter discussed various avenues of future research and how future research and development – combined with the efforts of this thesis – can be transitioned into a decision support tool. We discussed that an optimization-based decision support tool does not need to be an entirely new piece of software. In fact, a decision support tool should be designed to exist within the current software tools used by channel route planners. Furthermore, augmented crew composites, prepositioned crew composites, aerial refueling composites, and ramp transfer composites should be explored. Additionally, we suggest that hierarchical decomposition techniques and a rolling horizon approach should be explored and included in an optimization-based decision support tool.

[This Page Intentionally Left Blank]

7 Summary and Future Research Beyond Channel Routes

The focus of this research is to illustrate the ability of an optimization-based approach to design Air Mobility Command's (AMC) channel route network. Traditional network design formulations yield weak LP-bounds, undesirable solution times, and fail to capture difficult operational nuances. To overcome these shortfalls, we use a composite variable formulation, which can successfully handle complex operational rules and regulations, can be solved quickly, and can be easily manipulated to allow a wide range of sensitivity analysis without destroying the desired structure of the model. This chapter provides a summary of the work presented in this thesis, and we suggest several avenues of future research beyond channel routes.

7.1 Thesis Summary

In **Chapter 1**, we present the motivation for this research. We describe the Defense Transportation System (DTS) and AMC's role within the DTS. The air mobility environment changes continually, making it difficult for planners at the Tanker Airlift Control Center (TACC) to allocate flying hour requirements among AMC's four different mission areas – Special Assignment Airlift Missions (SAAMs), exercises, contingencies, and channel routes. The channel route network, which is used to transport military personnel and cargo throughout the

world, often suffers continual changes and cancellations as a result of channel route missions' low priority relative to the other mission areas, and the unpredictable nature of the higher priority mission areas. When channel route missions are changed or cancelled, either customer service is degraded or AMC is forced to pay commercial carriers premium rates to transport cargo throughout the world. Thus, a channel route network that is dynamically responsive is desired.

In **Chapter 2**, we provide an overview of AMC, laying the necessary groundwork to understand the context of the channel route network. We define the scope of this research to be the creation of an initial cut (i.e., a monthly schedule) for the portion of the channel route network used to transport cargo. In this chapter, we describe both the physical system and the current channel route planning and execution process. The physical system consists of military and commercial aircraft and aircrews, aerial ports (i.e., airports) located throughout the world, and cargo items that must flow between the aerial ports via the available set of aircraft. The current channel route planning process is largely manual, taking up to two months to generate a monthly schedule. Historically, automated methods at AMC have focused on wartime planning, leaving day-to-day operations to be scheduled manually. In this chapter, we also introduce AMC's three operational objectives: readiness, customer service, and net operating result (NOR). On a day-to-day level, there is no consideration of the system-wide impact that changes to the channel route schedule have on these operational objectives.

In **Chapter 3**, we develop two traditional network design formulations. The arc flow formulation (CRM-A) is intractable for realistic instances of the channel route network. We overcome the tractability issues of CRM-A using a path flow formulation (CRM-P) in conjunction with a decomposition approach based on classic Dantzig-Wolfe decomposition. In the decomposition, we solve a shortest path sub-problem for each aircraft type and each cargo commodity. The role of the sub-problems is to identify cost improving aircraft routes and cargo paths to be considered by the master level scheduler. We implement CRM-P and the decomposition using a small data set that consists of one week's cargo demands, and thirty aerial ports located throughout the eastern US, Europe, and the Middle-East. The results illustrate CRM-P's weak LP-relaxation, which is a common downfall of traditional network design formulations that are used to model large-scale transportation systems. The weak LP-bounds result from two sets of constraints: (1) forcing constraints, which cause fractional aircraft usage in the solution to the LP-relaxation, and (2) balance of flow constraints, which propagate fractional aircraft throughout the network.

Using the work of Armacost [3] and Armacost et al. [4], we formulate the monthly channel route planning problem using a composite variable formulation (CRM-C) in **Chapter 4**.

We remove the cargo flow decision variables and implicitly capture them within the design variables, which are the building blocks of composite variables. A composite variable is a single decision variable that corresponds to the selection of one or more aircraft missions that completely cover a subset of the cargo commodities. Creating composite variables requires more effort than is required to create variables for traditional network design formulations, but the benefit gained is significant. Using composite variables allows the channel route problem to be formulated without forcing constraints. Rather, we use a set of cover constraints, which are computationally superior. Composite variable formulations have tighter LP-relaxations and, consequently, much shorter solution times. We implement CRM-C using a small data set that consists of the same physical assets (i.e., aerial ports and aircraft) contained in the data set we use in Chapter 3. As a result of CRM-C's improved computational behavior, however, we are able to increase the size of the data set to span a 30-day planning horizon. The results illustrate the improved computational behavior of CRM-C.

In Chapter 5, we present results and analysis of both the computational characteristics of CRM-C (i.e., tractability and strength of the LP-relaxation) and the characteristics of the solution to CRM-C. We demonstrate that the data contained in a composite variable can be changed in a variety of ways without compromising the structural qualities of the model that result in tight LP-relaxations. Specifically, the composite variable generation procedures can be modified to change the nature of the problem, depending on the desired output. We illustrate this quality of CRM-C by testing various objective functions and various levels of the input parameters to the composite variable generation procedures. Each objective function is mapped to one of the three operational objectives of AMC or to the current planning process. We use a set of metrics to evaluate the characteristics (i.e., quality) of the solution. Similar to the objective functions, each metric captures one of AMC's three operational objectives or the current planning process. The results illustrate that the input parameters play a key role in defining the size of CRM-C. Furthermore, the objective function has a significant influence on the solution time and the quality of the LP-relaxation. Additionally, the results illustrate that under the four objective functions and metrics we investigate, the resulting solution performs well with respect to at most two of AMC's operational objectives. Although we do not conclude which solution is the best, we demonstrate that using a composite variable formulation allows numerous alternatives to be evaluated in much less time than currently required to generate a single initial cut.

In Chapter 6, we suggest several avenues of future research to improve CRM-C and to incorporate this research into an optimization-based decision support tool. We discuss integrating this research within existing software tools that are used by channel route planners,

and how such a tool would influence the current planning process. Additionally, we highlight several enhancements to CRM-C, including a hierarchical decomposition of the planning process, augmented crew composite variables, prepositioned crew composite variables, and ramp transfer composite variables. We outline the range of possible scenarios in which an optimization-based decision support tool would be useful.

In conclusion, the purpose of this thesis is to illustrate the ability of an optimization-based model to aid channel route planners in the design the channel route network. The channel route network is a large, complex system. Currently, channel route planners must use intuition and experience to design the channel route network. Unfortunately, the system is large enough that intuition often fails to grasp the system-wide impact of locally-made decisions. This research has provided the first steps toward an automated, optimization-based planning approach for the design of the channel route network. At this point, it is critical that real-world channel route planners evaluate the composite variable generation procedures and the solution to CRM-C. Incorporating the expertise of real-world channel route planners into the continued development of CRM-C is a rich opportunity to enhance this research.

7.2 Future Research

In Chapter 6, we highlight several avenues of future research in the context of CRM-C. In this section, we provide recommendations for future research in other air mobility arenas at AMC.

FLEET MIX

Aircraft acquisitions are costly for the US Air Force. The use of composite variable formulations to help analysts determine the optimal mix of aircraft should be explored. Aircraft with different capabilities could be easily incorporated into the composite variable generation procedures. As we illustrate in Chapter 5, various objective functions could be tested, and various metrics could be used to evaluate different fleet compositions.

AIRCRAFT ALLOCATION TO MISSION AREAS – “BARRELMASTER DECOMPOSITION”

The role of the barrelmaster is to allocate aircraft among the different mission areas. An interesting decomposition of the aircraft allocation problem faced by the barrelmaster is depicted in *Figure 7-1*. This approach is based on Dantzig-Wolfe decomposition. At the highest level, the barrelmaster is responsible for ensuring that flying hour requirements are satisfied for each aircraft type and that the number of military aircraft assigned does not exceed the number available. Additionally, the barrelmaster sends pricing information about each

aircraft type to lower level *mission area schedulers*. Each mission area scheduler determines the number of each aircraft type required to service the missions in their respective planning scope. The information from each mission area scheduler would then be combined to form a single composite variable to be considered by the barrelmaster. By construction, each composite mission variable would completely cover the missions from all four mission areas, alleviating the need for the barrelmaster to explicitly enforce these constraints. This does, however, require that the mission area schedulers have access to other airlift assets (i.e., commercial aircraft) because it is unlikely that the available number of military resources will completely cover all missions. The pricing information obtained from the barrelmaster would be used by each mission area scheduler to determine if assigning a commercial aircraft is less expensive.

In this decomposition, a composite variable corresponds to an entire allocation of aircraft among the mission areas. This is similar to Cohn's [22] use of composite variables to formulate a service parts logistics problem. Specifically, she uses a composite variable to capture the assignment of an entire set of warehouses to customers so that all customers are covered by a warehouse, which alleviates the need to explicitly state warehouse-customer cover constraints in the mathematical programming formulation. In the same sense, a single composite variable in the decomposition depicted in *Figure 7-1* ensures that each mission from the four mission areas is covered. Column generation plays a pivotal role in this decomposition. Consequently, column generation techniques for composite variables (see Cohn [22]) need to be explored.

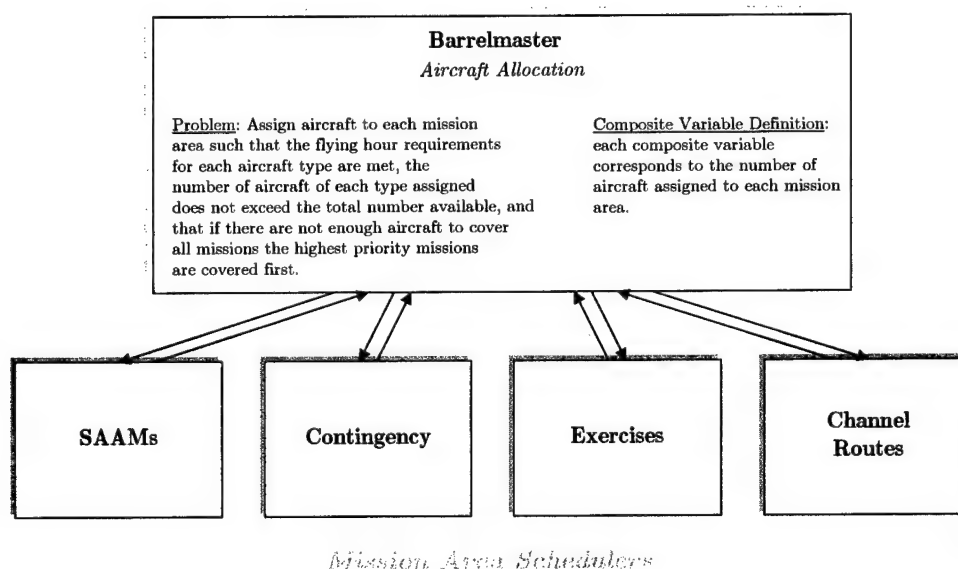


Figure 7-1: Proposed Barrelmaster Decomposition

AEROMEDICAL ROUTE STRUCTURE

Another interesting application for composite variable modeling is the CONUS aeromedical route structure. The US Air Force moves medical patients throughout the network of CONUS military bases on a daily basis, using a single aircraft type. The problem is different from designing the channel route network in the sense that the demands (i.e., the patients) in the aeromedical problem are extremely small compared to the aircraft capacity. Furthermore, the demands are not known until days, if not hours, before the mission must be flown. Given the patient movements, weather, aircraft availability, and base operating hours for a particular day, composite variables could be used to design the set of flight sequences for the day.

WARTIME SCENARIO ANALYSIS

Analysts at AMC are constantly evaluating "what-if" scenarios that require air mobility assets. We suggest that composite variable modeling for these scenarios be explored. Specifically, composite variable formulations would be helpful to determine the air mobility route structure for buildup and sustainment operations. Composite variable modeling would be a good fit with this type of analysis because composite variable formulations can be solved quickly, and they are extremely flexible. Analysts would be able to easily change the nature of the problem – without destroying the model's structure – by modifying the composite variable generation procedures. Furthermore, experience demonstrates that composite variable modeling successfully handles complex operational rules and regulations, which are inherent in the mobility operations considered in "what-if" scenario planning.

MODELING THE LOGISTICS PIPELINE

We suggest composite variable modeling techniques for modeling the entire military logistics pipeline be explored. In this context, a single composite variable would correspond to various modes of travel (i.e., land, air, sea) that a cargo item would use as it is shipped from its origin to its destination. For example, consider two cargo commodities: cargo commodity A (manufactured at a facility in Kansas and must be delivered to a customer in Europe) and cargo commodity B (manufactured in Tennessee and must be delivered to the same customer in Europe). A single composite variable that covers these cargo items would consist of a train carrying cargo commodity A from Kansas to an east coast aerial port, a truck delivering cargo commodity B from Tennessee to the same east coast aerial port, an aircraft carrying both cargo commodities to an aerial port in Europe, and a truck transporting both cargo commodities from the European aerial port to the customer's destination. We illustrate this composite variable in *Figure 7-2*. Modeling the logistics pipeline in this fashion would require the use of column generation techniques to deal with the large number of variables, which could potentially be in




Figure 7-2: A composite variable for the logistics pipeline problem

the billions. An important aspect of composite variable modeling is how composite variables are defined – that is, what decisions are combined (see Cohn [22]). *Figure 7-2* illustrates a composite variable that is created as a result of combining decisions about different modes of transportation. Another approach would be to combine the decisions corresponding to a single mode of transportation and leverage the fact that the cargo transfer points (i.e., the locations cargo changes modes) are known *a priori*.

[This page intentionally left blank]

Appendix A: Formulations

$$\begin{aligned}
 \text{MCNF-A} = \min \quad & \sum_{k \in K} \sum_{(i,j) \in A} c_{ij}^k x_{ij}^k \\
 \text{s.t.} \quad & \sum_{k \in K} x_{ij}^k \leq u_{ij} \quad \forall (i,j) \in A, \\
 & \sum_{\{j:(i,j) \in A\}} x_{ij}^k - \sum_{\{j:(j,i) \in A\}} x_{ji}^k = \begin{cases} b^k, & \text{if } i = O(k) \\ -b^k, & \text{if } i = D(k) \\ 0, & \text{otherwise} \end{cases} \quad \forall i \in N, k \in K, \\
 & x_{ij}^k \geq 0 \quad \forall (i,j) \in A, k \in K.
 \end{aligned}$$

$$\begin{aligned}
 \text{MCNF-P} = \min \quad & \sum_{k \in K} \sum_{p \in P^k} c_p^k x_p^k \\
 \text{s.t.} \quad & \sum_{k \in K} \sum_{p \in P^k} \delta_{ij}^p x_p^k \leq u_{ij} \quad \forall (i,j) \in A, \\
 & \sum_{p \in P^k} x_p^k = b^k \quad \forall k \in K, \\
 & x_p^k \geq 0 \quad \forall k \in K, p \in P^k.
 \end{aligned}$$

$$\begin{aligned}
 \text{NDF} = \min \quad & \sum_{k \in K} \sum_{p \in P^k} c_p^k x_p^k + \sum_{f \in F} \sum_{(i,j) \in A} d_{ij}^f y_{ij}^f \\
 \text{s.t.} \quad & \sum_{k \in K} \sum_{p \in P^k} \delta_{ij}^p b^k x_p^k \leq \sum_{f \in F} u_{ij}^f y_{ij}^f \quad \forall (i,j) \in A, \\
 & \sum_{p \in P^k} x_p^k = 1 \quad \forall k \in K, \\
 & \sum_{ij} y_{ij}^f = 0 \quad \forall f \in F, \\
 & x_p^k \geq 0 \quad \forall k \in K, p \in P^k, \\
 & y_{ij}^f \in \{0,1\} \quad \forall (i,j) \in A, f \in F.
 \end{aligned}$$

CRM-A =

$$\begin{aligned}
& \min \sum_{f \in F} \sum_{a \in \bar{A}[t^c]} y_a^f + \sum_{f \in F} \sum_{a \in \underline{A}[t^c]} g_a^f \\
& s.t. \quad \sum_{k \in K} x_a^k \leq u_a^{FB} + \sum_{f \in F} u^f y_a^f \quad \forall a \in \bar{A}, \\
& \quad \quad \quad \mathcal{N}x^k = b^k \quad \forall n \in N, k \in K, \\
& \quad \quad \quad \sum_{a \in \underline{A}[t^c]} g_a^f + \sum_{a \in \bar{A}[t^c]} y_a^f \leq \rho^f \quad \forall f \in F, \\
& \quad \quad \quad \mathcal{D}y^f = 0 \quad \forall f \in F, \\
& \quad \quad \quad \sum_{f \in F^m} \sum_{a \in \bar{A}} y_a^f h_a^f \leq \bar{H}^m \quad \forall m \in M, \\
& \quad \quad \quad \sum_{f \in F^m} \sum_{a \in \bar{A}} y_a^f h_a^f \geq \underline{H}^m \quad \forall m \in M, \\
& \quad \quad \quad \sum_{f \in F} \sum_{a \in \bar{A}[n^-]} y_a^f \leq e_n \quad \forall n \in N^{B^s}, \\
& \quad \quad \quad \sum_{f \in F} g_a^f \leq e_n \quad \forall n \in N^{B^{NS}}, a \in \underline{A} : O(a) = B(n) \\
& \quad \quad \quad y_a^f = 0 \quad \forall a \in \bar{A} : t \in T^{ops}, f \in F, \\
& \quad \quad \quad x_a^k \geq 0 \quad \forall k \in K, a \in A, \\
& \quad \quad \quad y_a^f, g_a^f \in \mathbb{Z}^+ \quad \forall f \in F, a \in A.
\end{aligned}$$

CRM-P =

$$\begin{aligned}
& \min \sum_{f \in F} \sum_{r \in R^f} \sum_{a \in A[l^c]} \delta_a^r y_r^f \\
& \text{s.t.} \quad \sum_{k \in K} \sum_{p \in P^k} \delta_a^p b^k x_p^k \leq u_a^{FB} + \sum_{f \in F} \sum_{r \in R^f} \delta_a^r u^f y_r^f \quad \forall a \in \bar{A}, \\
& \quad \sum_{p \in P^k} x_p^k = 1 \quad \forall k \in K, \\
& \quad \sum_{r \in R^f} \sum_{a \in A[l^c]} \delta_a^r y_r^f \leq \rho^f \quad \forall f \in F, \\
& \quad \Phi y^f = 0 \quad \forall f \in F, \\
& \quad \sum_{f \in F^m} \sum_{r \in R^f} \sum_{a \in \bar{A}} \delta_a^r h_a^f y_r^f \leq \bar{H}^m \quad \forall m \in M, \\
& \quad \sum_{f \in F^m} \sum_{r \in R^f} \sum_{a \in \bar{A}} \delta_a^r h_a^f y_r^f \geq \underline{H}^m \quad \forall m \in M, \\
& \quad \sum_{f \in F} \sum_{r \in R^f} \sum_{a \in \bar{A}} \delta_a^r y_r^f \leq e_n \quad \forall n \in N^{B^s}, \\
& \quad \sum_{f \in F} \sum_{r \in R^f} \sum_{a \in \bar{A}} \delta_a^r y_r^f \leq e_n \quad \forall n \in N^{B^{NS}}, \\
& \quad \sum_{f \in F} \sum_{r \in R^f} \delta_w^r y_r^f + \psi_w^{FB} \geq \xi^w \quad \forall w \in W, \\
& \quad x_p^k \geq 0 \quad \forall k \in K, p \in P^k, \\
& \quad y_r^f \in \mathbb{Z}^+ \quad \forall f \in F, r \in R^f.
\end{aligned}$$

CRM-C =

$$\begin{aligned}
 \min \quad & \sum_{c \in \mathbb{C}^O} \sum_{f \in F} \eta_c^f z_c + \sum_{n \in N^{gl(rul)}} \sum_{f \in F} g_{n^-}^f \\
 \text{s.t.} \quad & \sum_{c \in \mathbb{C}^O} \delta_c^k z_c + \sum_{c \in \mathbb{C}^R} \delta_c^k z_c \geq 1 \quad \forall k \in K, \\
 & \sum_{c \in \mathbb{C}^O} \beta_c^w z_c + \sum_{c \in \mathbb{C}^R} \beta_c^w z_c \geq \xi^w \quad \forall w \in W, \\
 & \sum_{c \in \mathbb{C}^O} \phi_c^f(n^+) z_c + g_{n^+}^f - \sum_{c \in \mathbb{C}^O} \phi_c^f(n^-) z_c - g_{n^-}^f = 0 \quad \forall n \in N^{D^s}, f \in F^n, \\
 & \sum_{c \in \mathbb{C}^R} \gamma_c^l z_c = 1 \quad \forall l \in FB, \\
 & \sum_{c \in \mathbb{C}^O} \eta_c^f z_c + \sum_{n \in N^{gl(rul)}} g_{n^-}^f \leq \rho^f \quad \forall f \in F, \\
 & \sum_{c \in \mathbb{C}^O} \tau_c^m z_c \leq \overline{H}^m \quad \forall m \in M, \\
 & \sum_{c \in \mathbb{C}^O} \tau_c^m z_c \geq \underline{H}^m \quad \forall m \in M, \\
 & g_{n^+}^f \geq 0 \quad \forall n \in N, f \in F, \\
 & g_{n^-}^f \geq 0 \quad \forall n \in N, f \in F, \\
 & z_c \in \mathbb{Z}^+ \quad \forall c \in \mathbb{C}.
 \end{aligned}$$

Appendix B: Glossary of Acronyms

AB	Air Base
ACC	Air Combat Command
ADANS	Airlift Deployment Analysis System
AFB	Air Force Base
AFM	Airlift Flow Model
AMC	Air Mobility Command
APOD	Aerial Port of Debarkation
APOE	Aerial Port of Embarkation
ARA	Aircraft Readiness Account
ARC	United States Air Force Reserve and National Guard Component
CAMPS	Consolidated Air Mobility Planning System
CDD	Crew Duty Day
CONUS	Continental United States
CRAF	Civil Reserve Air Fleet
DeCA	Defense Commissary Agency
DIPs	Diplomatic clearances
DLA	Defense Logistics Agency
DoD	Department of Defense
DOS	Department of State
DTS	Defense Transportation Network
FBI	Federal Bureau of Investigation
FHP	Flying Hour Program
GDSS	Global Decision Support System
IP	Integer Program
JA/ATT	Joint Airborne/Air Transportability Training
JCS	Joint Chiefs of Staff
MAC	Military Airlift Command
MASS	Mobility Analysis Support System
MIP	Mixed Integer Program
MOG	Maximum on Ground
MSC	Military Sealift Command
MTMC	Military Traffic Management Command

OCONUS	Outside the Continental United States
PMOG	Parking Maximum on Ground
SAAM	Special Assignment Airlift Mission
SAC	Strategic Air Command
TAC	Tactical Air Command
TACC	Tanker Airlift Control Center
TCC	Transportation Component Command
TCN	Transportation Control Number
TWCF	Transportation Working Capital
US	United States
USA	United States Army
USN	United States Navy
USAF	United States Air Force
USTRANSCOM	United States Transportation Command
WMOG	Working Maximum on Ground

References

- [1] Ackley, M., B.Carter, J.Litko, K.Ware, and A.Whisman. "Decision Support for the Air Mobility Command Scheduled Cargo System." Air Mobility Command Working Paper, May 1993.
- [2] Ahuja, R., T. Magnanti, and J. Orlin. *Network Flows: Theory, Algorithms, and Applications*. Prentice Hall, Upper Saddle River, NJ, 1993.
- [3] Armacost, A. "Composite Variable Formulations for Express Shipment Service Network Design." PhD Thesis in Operations Research. Massachusetts Institute of Technology, Cambridge, MA, September 2000.
- [4] Armacost, A., C. Barnhart, and K. Ware. "Composite Variable Formulations for Express Shipment Service Network Design." *Transportation Science*, Vol. 36, No. 1, pp.1-20, February 2002.
- [5] Baker, Steven F., D. Morten, R. Rosenthal, L.Williams. "Optimizing Military Airlift." Report, July 2000.
- [6] Barnhart, C., N. Boland, L. Clarke, E. Johnson, G. Nemhauser, and R. Shenoi. "Flight String Models for Aircraft Fleeting and Routing." *Transportation Science*, Vol. 32, No. 3, pp. 208-220.
- [7] Barnhart, C., A. Farahat, M. Lohatepanont. "Airline Fleet Assignment: An Alternative Model and Solution Approach based on Composite Variables." Submitted to *Operations Research*.
- [8] Barnhart, C., E. Johnson, G. Nemhauser, M. Savelsbergh, and P. Vance. "Branch-and-Price: Column Generation for Huge Integer Programs." *Operations Research*, Vol. 46, No. 3, pp. 316-329, May-June 1998.
- [9] Barnhart, C., E. Johnson, G. Nemhauser, and P. Vance. "Crew Scheduling." *Handbook of Transportation Science*. Ed. Randolph W. Hall. Kluwer Academic Publisher, Norwell, MA, 1999. pp. 493-521
- [10] Barnhart, C., F. Lu, and R. Shenoi. "Integrated Airline Schedule Planning." *Operations Research in the Air Industry; International Series in Operations Research and Management Science*. Vol. 9. Ed. Gang Yu, Kluwer Academic Publisher, Norwell, MA, 1999. pp. 384-403
- [11] Barnhart, C., and R. Schneur. "Air Network Design For Express Shipment Service." *Operations Research*, Vol.44, pp.852-863, November-December 1996.

- [12] Barnhart, C., and K. Talluri. "Airline Operations Research." *Design and Operations of Civil and Environmental Engineering Systems*, pp. 435-469, 1997.
- [13] Barth, C., "Composite Mission Variable Formulation for Real-Time Mission Planning." Masters Thesis in Operations Research. Massachusetts Institute of Technology, Cambridge, MA, June 2001.
- [14] Becker, M., and S. Smith. "Mixed-Initiative Resource Management: The AMC Barrel Allocator." *Proceedings 5th International Conference on Artificial Intelligence Planning and Scheduling (AIPS-2000)*, Breckenridge, CO, April, 2000.
- [15] Bertsimas, D., and J.N. Tsitsiklis. *Introduction to Linear Optimization*. Athena Scientific, Belmont, MA, 1998.
- [16] Brandt, C.M, C.J. Burke, K.W. Currie, A.R. Heminger, T.L. Pohlen, and D.K. Vaughn. "Dynamic Response Logistics: Changing Environments, Technologies, And Processes." *Air Force 2025*. Vol. 2, Chpt. 2, Air University, 1996.
- [17] Büdenbender, K., T. Grunert, and H-J. Sebastian. "A Hybrid Search/Branch-and-Bound Algorithm for the Direct Flight Network Design Problem." *Transportation Science*, Vol. 34, No. 4, pp. 364-380. November 2000.
- [18] Busch, I., G. Harrison, M.R. Hilliard, R.D. Kraemer, C.Liu, R.S. Solanki. "Scheduling the Operation Desert Storm Airlift: An Advanced Automated Scheduling Support System." *Interfaces*, Vol. 22, No. 1, pp.131-146, January-February 1992.
- [19] Bowlin, W.F. "Assessing the Financial Condition of Civil Reserve Air Fleet Participants." *Military Operations Research*, Vol. 6, No. 1, pp.43-56, 2001.
- [20] Clark, S.J., "Large-Scale Optimization Planning Methods for Distribution and Transportation Logistics: An Application to Contingency Munitions." Masters Thesis in Operations Research. Massachusetts Institute of Technology, Cambridge, MA, June 2001.
- [21] Cochard, D. and K. Yost. "Improving Utilization of Air Force Cargo Aircraft." *Interfaces*, Vol. 15, No. 1, pp.53-68, January-February 1985.
- [22] Cohn, A., "Composite-Variable Modeling for Large-Scale Problems in Transportation and Logistics." PhD Thesis in Operations Research. Massachusetts Institute of Technology, Cambridge, MA, June 2002.
- [23] Cohn, A., and C. Barnhart. "Improving Crew Scheduling by Incorporating Key Maintenance Routing Decisions." Submitted to *Operations Research*, December 2000.
- [24] Cordeau, J., G. Stojkovic, F. Soumis, and J. Desrosiers. "Benders Decomposition for Simultaneous Aircraft Routing and Crew Scheduling." *Technical Report G-2000-37*, GERAD, École Polytechnique de Montréal. 2000.

- [25] Crainic, T., "Service network design in freight transportation." *European Journal of Operations Research*, Vol. 122, pp. 272-288, 2000.
- [26] Desrochers, M., and F. Soumis. "A Generalized Permanent Labeling Algorithm for the Shortest Path Problem with Time Windows." *INFOR*. Vol. 26, pp.191-212, 1988.
- [27] Dragoo, G.F., J. Edgell, L.W. Jackson, and S.K. Spangler. "Logistics in 2025: Consider It Done." *Air Force 2025*. Vol. 2, Chpt. 1, Air University, 1996.
- [28] Farkas, A. "The Influence of Network Effects and Yield Management on Airline Fleet Assignment Decisions." PhD Thesis. Massachusetts Institute of Technology, Cambridge, MA, 1996.
- [29] Gendron, B., T. Crainic, A. Frangioni. "Multicommodity capacitated network design." Eds. B. Sanso and P. Soriano, Kluwer Academic Publishers, Dordrecht, The Netherlands, 1999. pp. 1-19
- [30] Golden, B.L., D.S. Feshbach, L.S. Levy, and H.K. Rappoport. "Estimating Loads of Aircraft for the Military Airlift Command." *Interfaces*, Vol. 21, No. 4, pp.63-78, July-August 1991.
- [31] Golden, B.L., L.S. Levy, H.K. Rappoport, and K.J. Toussaint. "A Planning Heuristic for Military Airlift." *Interfaces*, Vol. 22, No.3, pp.73-87, May-June 1992.
- [32] Golden, B.L., L.S. Levy, H.K. Rappoport, and K.J. Toussaint. "A transportation problem formulation for the MAC Airlift Planning Problem." *Annals of Operations Research*, Vol.50, pp.505-523, 1994.
- [33] Grünert, T., and H-J. Sebastian. "Planning Models for long-haul operations of postal and express shipment companies." *European Journal of Operational Research*, Vol. 122, No. 2, pp.289-309, 2000.
- [34] Hane, C., C. Barnhart, E. Johnson, R. Marsten, G. Nemhauser, and G. Sigismondi. "The Fleet Assignment Problem: Solving a Large-Scale Integer Program." *Mathematical Programming*, Vol. 70, No. 2, pp. 211-232, 1995.
- [35] Hughes, G.C., and R. Roehrkasse. "Crisis Analysis: Operation Desert Shield." *OR/MS Today*, pp.22-27, December 1990.
- [36] *Joint Vision 2010*. Joint Chiefs of Staff, U.S. Government Printing Office, Washington DC. 1997.
- [37] *Joint Vision 2020*. Joint Chiefs of Staff, U.S. Government Printing Office, Washington DC. 2000.
- [38] Kim, D., C. Barnhart, G. Reinhardt, and K. Ware. "Multimodal Express Package Delivery: A Service Network Design Application." *Transportation Science*, Vol. 33, No. 4, pp.391-407, November 1999.

- [39] Klabjan, D., E. Johnson, G. Nemhauser, E. Gelman, and S. Ramaswamy. "Airline Crew Scheduling with Time Windows and Plane Count Constraints." Technical Report TLI/LEC-99-12, Georgia Institute of Technology, 1999.
- [40] Liu, C., and R. Solanki. "The Allocation of Cargo to Channel Missions." Presented at ORSA/TIMS National Meeting, 1993.
- [41] Luebbbers, D., "FY02 Fixed Buy Program: Channel Cargo and Passenger Routes." Presentation, TACC/XOG, 2001.
- [42] Magnanti, T., and R. Wong. "Network Design and Transportation Planning: Models and Algorithms." *Transportation Science*, Vol. 18, No. 1, pp 1-49, February 1984.
- [43] McDaniel, Janet. GS12, TACC/XOG. *Personal Interview*, 20 August 2001.
- [44] Nemhauser, G., and L. Wolsey. *Integer and Combinatorial Optimization*. John Wiley & Sons, Inc., New York, NY, 1999.
- [45] Rexing, B., C. Barnhart, A. Jarrah, T. Kniker, and N. Krishnamurthy. "Airline Fleet Assignment with Time Windows." *Transportation Science*, Vol. 34, No. 1, pp. 1-20, 2000.
- [46] Schrader, C.R., "Gulf War Logistics." *Parameters*, Vol. 25, No. 4, pp. 142-48, Winter, 1995-1996.
- [47] Solanki, R.S., and F. Southworth. "An Execution Planning Algorithm for Military Airlift." *Interfaces*, Vol. 21, No. 4, pp.-121-131, July-August 1991.
- [48] United States Air Mobility Command Report. Stakeholders Report: Air Mobility Command, Fiscal Year 2000. 2001.
- [49] United States Air Mobility Command Website. *Brief History of the Air Mobility Command*. http://public.scott.af.mil/hqamc/library/amc_history.htm, 18 February 2002.
- [50] United States Air Mobility Command, Tanker Airlift Control Center. *TACC Informational Brochure*. June 2000.
- [51] United States Transportation Command Handbook 24-2. *Understanding the Defense Transportation System, 3rd Edition*. September 2000.
- [52] United States Transportation Command Report. Transportation For A New Millennium, 2000 Annual Command Report. 2001.
- [53] Vance, P., C. Barnhart, E. Johnson, and G. Nemhauser. "Airline Crew Scheduling: A New Formulation and Decomposition Algorithm." *Operations Research*, Vol. 45, No. 2, pp. 188-200, March-April 1997.

- [54] Wright, S.A., "Covalidation of Dissimilarly Structured Models." PhD Dissertation in Operations Research. Air Force Institute of Technology, Wright-Patterson Air Force Base, OH, March 2001.

[This page intentionally left blank]